
PHENOMENOLOGICAL STUDIES
ON SUPERSYMMETRY
AND THE STRONG FORCE

PETER Z. SKANDS

DEPARTMENT OF THEORETICAL PHYSICS
LUND UNIVERSITY, SWEDEN

DOCTORAL THESIS

FOR THE DEGREE OF DOCTOR OF PHILOSOPHY IN PHYSICS

THESIS ADVISOR:

TORBJÖRN SJÖSTRAND (LUND)

FACULTY OPPONENT:

MICHELANGELO MANGANO (CERN)

WITH THE PERMISSION OF THE FACULTY OF NATURAL SCIENCES OF LUND UNIVERSITY.
TO BE PRESENTED FOR PUBLIC CRITICISM IN LECTURE HALL F OF THE DEPARTMENT OF
THEORETICAL PHYSICS ON FRIDAY OCTOBER 8TH, 2004, AT 13.30.

Organization LUND UNIVERSITY Department of Theoretical Physics Sölvegatan 14A SE-223 62 Lund		Document Name DOCTORAL DISSERTATION	
		Date of issue September 2004	
		CODEN:	
Author(s) Peter Zeiler Skands		Sponsoring organization	
Title and subtitle Phenomenological Studies on Supersymmetry and the Strong Force			
Abstract In this thesis, we first consider the phenomenology of R -parity violating supersymmetric extensions of the Standard Model. Specifically, a large number of Lepton and Baryon Number violating decays of sparticles to particles are implemented in the PYTHIA event generator. The augmented generator has then been used to study event properties of these models, mainly at the LHC. During these studies, the development of the so-called SuSy Les Houches Accord opened for an easier interfacing between SuSy calculational tools. A first 'proof-of-concept' application of the Accord was to input SuSy spectra and decay tables calculated by the SPHENO program into PYTHIA, for the purpose of determining whether experiments at the LHC will be sensitive to a connection between neutrino mixing and SuSy with bilinear Lepton Number violation. For the case of Baryon Number violation, an extension of the Lund String Model to colour topologies involving explicit Baryon Numbers is developed, called 'junction fragmentation', which enables us to trace the migration of the generated Baryon number in some detail. This model is then itself extended and applied to improve the description of the physics of beam remnants in hadron-hadron collisions, in the context of a new framework for minimum-bias and underlying events. Based on the idea of multiple perturbative parton-parton interactions in hadron collisions, a sophisticated treatment of correlations in impact parameter, flavour, momentum, and colour is introduced. The last step so far in the evolution of this model has been the development of p_{\perp} -ordered initial- and final-state radiation algorithms, which among other things have enabled us to 'interleave' initial-state radiation and multiple interactions into one common sequence of decreasing p_{\perp} values. Nonetheless, there is still a long way to go before a clear and unique picture of hadron collisions may emerge.			
Key words QCD, parton showers, hadronisation, string fragmentation, hadron collisions, underlying event, beam remnants, SuSy, R -parity violation.			
Classification system and/or index terms (if any)			
Supplementary bibliographical information		Language English	
ISSN and key title		ISBN 91-628-6189-1	
Recipient's notes		Number of pages 324	Price
		Security classification	

Distribution by (name and address)

Peter Skands, Dept. of Theoretical Physics, Sölvegatan 14A, SE-223 62 Lund

I, the undersigned, being the copyright owner of the abstract of the above-mentioned dissertation, hereby grant to all reference sources the permission to publish and disseminate the abstract of the above-mentioned dissertation.

Signature _____

Date 25 Aug. 2004 _____

*Hun synger om vår
om daggry i maj
solstråler danser
på gyldenblond hår*

*Mens engene vågner
og morgenen gryer
leger sylfider
blandt tågernes slør*

*I et land bort om lande
bort om ord, bort om drømme
i elvernes dal
ved nymfernes vande*

*Forsvinder nu duggen
for morgenens glød
og dawning blir dag
i min elskedes favn.*

This thesis is based on the following publications:

- I **Searching for L-Violating Supersymmetry at the LHC.**
By P. Skands.
LU TP 01-32, Oct 2001.
Published in European Physical Journal **C23** (2002) 173.
- II **Baryon Number Violation and String Topologies.**
By T. Sjöstrand and P. Skands.
LU TP 02-46, Dec 2002.
Published in Nuclear Physics **B 659** (2003) 243.
- III **SUSY Les Houches Accord: Interfacing SUSY Spectrum Calculators, Decay Packages, and Event Generators.**
By P. Skands, B.C. Allanach, H. Baer, C. Balázs, G. Bélanger, F. Boudjema, A. Djouadi, R. Godbole, J. Guasch, S. Heinemeyer, W. Kilian, J-L. Kneur, S. Kraml, F. Moortgat, S. Moretti, M. Mühlleitner, W. Porod, A. Pukhov, P. Richardson, S. Schumann, P. Slavich, M. Spira, G. Weiglein.
LU TP 03-39, Nov 2003.
Published in Journal of High Energy Physics **07** (2004) 036.
- IV **Measuring Neutrino Mixing angles at LHC.**
By W. Porod and P. Skands.
LU TP 03-50, ZU-TH 20/30, Jan 2004. [hep-ph/0401077]
In Beyond the Standard Model Working Group: Summary report, 3rd Les Houches Workshop: Physics at TeV Colliders, Les Houches, France, 26 May - 6 Jun 2003, B. C. Allanach *et al.* [hep-ph/0402295].
- V **Multiple Interactions and the Structure of Beam Remnants.**
By T. Sjöstrand and P. Skands.
LU TP 04-01, Feb 2004.
Published in Journal of High Energy Physics **03** (2004) 053.
- VI **Transverse-Momentum-Ordered Showers and Interleaved Multiple Interactions.**
By T. Sjöstrand and P. Skands.
LU TP 04-29, Aug 2004.
To be submitted to the European Physical Journal **C**.

During my Ph.D. studies, I have also written and contributed to several other articles. These are not reproduced in this thesis, but are listed here for completeness:

- VII **Baryon Number Violation and String Hadronization.**
By P. Skands.
LU TP 02-35, Sep 2002. [hep-ph/0209199]
In the proceedings of 10th International Conference on Supersymmetry and Unification of Fundamental Interactions (SUSY02), Hamburg, Germany, 17-23 Jun 2002, eds. P. Nath, P. M. Zerwas, and C. Grosche (DESY, Hamburg, 2002), vol. 2, p. 831.
- VIII **PYTHIA 6.3 Physics and Manual.**
By T. Sjöstrand, L. Lönnblad, S. Mrenna, P. Skands.
LU TP 03-38, Aug 2003. [hep-ph/0308153]
- IX **Progress on Multiple Interactions.**
By P. Skands and T. Sjöstrand.
LU TP 03-45, Oct 2003. [hep-ph/0310315]
In the proceedings of International Europhysics Conference on High-Energy Physics (HEP 2003), Aachen, Germany, 17-23 Jul 2003, European Physical Journal C Direct, DOI: 10.1140/epjcd/s2003-03-520-7.
- X **The SUSY Les Houches Accord Project**
By P. Skands, B.C. Allanach, H. Baer, C. Balázs, G. Bélanger, F. Boudjema, A. Djouadi, R. Godbole, J. Guasch, S. Heinemeyer, W. Kilian, J-L. Kneur, S. Kraml, F. Moortgat, S. Moretti, M. Mühlleitner, W. Porod, A. Pukhov, P. Richardson, S. Schumann, P. Slavich, M. Spira, G. Weiglein.
LU TP 03-48, Dec 2003.
In Beyond the Standard Model Working Group: Summary report, 3rd Les Houches Workshop: Physics at TeV Colliders, Les Houches, France, 26 May - 6 Jun 2003, B. C. Allanach *et al.* [hep-ph/0402295].
- XI **Multiple Interactions and Beam Remnants.**
By T. Sjöstrand and P. Skands.
LU TP 04-04, Jan 2004. [hep-ph/0401060]
In The QCD/SM working group: Summary report, 3rd Les Houches Workshop: Physics at TeV Colliders, Les Houches, France, 26 May - 6 Jun 2003, M. Dobbs *et al.* [hep-ph/0403100].

XII Les Houches Guidebook to Monte Carlo Generators for Hadron Collider Physics.

By M. A. Dobbs, S. Frixione, E. Laenen, K. Tollefson, H. Baer, E. Boos, B. Cox, R. Engel, W. Giele, J. Huston, S. Ilyin, B. Kersevan, F. Krauss, Y. Kurihara, L. Lönnblad, F. Maltoni, M. Mangano, S. Odaka, P. Richardson, A. Ryd, T. Sjöstrand, P. Skands, Z. Wqs, B. R. Webber, D. Zeppenfeld.

Mar 2004. [hep-ph/0403045]

In The QCD/SM working group: Summary report, 3rd Les Houches Workshop: Physics at TeV Colliders, Les Houches, France, 26 May - 6 Jun 2003, M. Dobbs *et al.* [hep-ph/0403100].

PHENOMENOLOGICAL STUDIES
ON SUPERSYMMETRY
AND THE STRONG FORCE

Peter Z. Skands

Contents

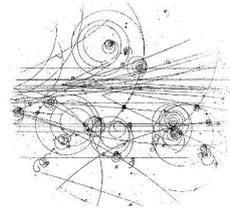
Introduction	1
I Searching for L-Violating Supersymmetry at the LHC	31
II Baryon Number Violation and String Topologies	57
III SUSY Les Houches Accord: Interfacing SUSY Spectrum Calculators, Decay Packages, and Event Generators	127
IV Measuring Neutrino Mixing angles at LHC	165
V Multiple Interactions and the Structure of Beam Remnants	173
VI Transverse-Momentum-Ordered Showers and Interleaved Multiple Interactions	255

Introduction

*In your hands now here's my book,
between whose pages, if you look,
is written verse on verse,
in language dry and terse;
about some tiny particles,
obeying mighty principles!
For instance supersymmetry,
or the $SU(3)$ of QCD.
Before you in this splendid mess,
there is reason nonetheless.
Sure, it may be just a guess,
but it's an educated one I stress.*

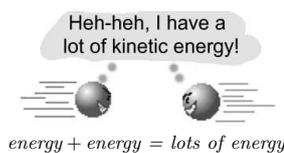
Imagine embarking on an expedition in the footsteps of the reputable 20th century explorer Mr. Tompkins [1], on a voyage down through the atom, inside the nucleus, to the quantum world of subatomic particles. Placing ordinary common sense under strict orders to keep quiet, we find ourselves in the natural habitat of elementary particles. Of these, a fortunate few have been observed in nature and enjoy the privilege of being considered fundamental. Many more lead a shadowy and more hypothetical life, never yet having been sighted in actual existence. In this thesis, both types will be encountered, from well-known quarks and gluons to speculative sparticles (spelling intended). This introduction then, is meant to acquaint the non-expert with the chimerical wonderland that they inhabit and that it has been my great privilege to visit frequently these last few years, so often in the best of company.

Essentially, all that is required to make this journey, from the fragile world of our immediate senses to the habitual haunts of Mr. Tompkins, is a perception tuned to rather small wavelengths, or equivalently, very high energies. The proper equipment for the expedition normally consists of, firstly, the *vehicle*, a good-sized accelerator, secondly, the *engine*, an amount of electricity that otherwise might keep a small city comfortable, and, thirdly, the *eye*, a hi-tech wonder composed of gas, steel, lead, cryogenics, computers, and electronics, the detector. Thus equipped, the tiny world of elementary particles, beginning at length scales



below a millionth of a nanometre, may be opened to human investigation. Theorists, of course, will undertake the same expedition without all the equipment and are therefore much cheaper, though somewhat less reliable.

Today, particle accelerators are capable of reaching energy scales several billion times larger than those of ordinary chemical reactions. For instance, the energy released by the combustion of an octane molecule in a car engine is about 50 electron-Volts¹, whereas the Tevatron at Fermilab currently produces proton-antiproton collisions at an energy of 1.960.000.000.000 electron-Volts, corresponding to each of the beams having been accelerated through a potential difference of some 980 billion Volts. At these energies, the ordinary phases of matter (solids, liquids, gasses, and even plasmas) cease to be meaningful. We are in a domain where the effects of relativity and quantum mechanics are highly noticeable; matter and antimatter can be created and annihilated from the vacuum, and particles are not well localised but are described by fuzzy quantum wave functions, or ‘fields’.



In fact, this fundamental “fuzziness” of Nature is one of the most striking aspects of quantum physics. More precisely, any conceivable measurement is performed at some intrinsic “resolution scale”. Observing visible light, for instance, we are perceiving waves in the electromagnetic field whose crests are separated by no less than several hundred nanometres. Due to the fuzziness of the light quanta, visible light could never tell us about structures smaller than that. That is why ‘particle physics’ and ‘high energy physics’ are synonymous; to resolve things that are, at the largest, 100 million times smaller than the distance between two crests of a wave of visible light, we need probes of extremely short wavelengths. That is precisely what is obtained by accelerating matter to high energies. The higher the energy, the better our “vision” in the fuzzy world of Mr. Tompkins.

Returning to antimatter, the discovery of the anti-electron, or *positron*, dates as far back as 1932. Today, antipartners have been found for all of the known matter particles. Within measurement errors, the antipartners have exactly the same masses as their particle counterparts, but they carry diametrically opposite charges. In the quantum theory of fields, this phenomenon has a quite elegant and natural explanation. Matter and antimatter are simply two different states of the *same* field, hence you can’t have one without the other. A more intuitive way of looking at antimatter is simply to note that in many respects, it acts just like matter travelling *backwards* through time.

Matter and force are also intimately related – by symmetries. Among the true success stories of 20th century theoretical physics was the discovery that matter fields can generate force fields in a quite natural way. Specifically, in

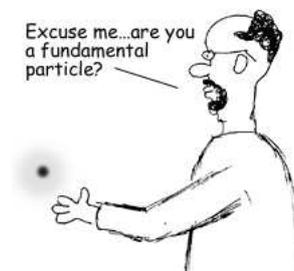
¹One electron-Volt is the energy gained by an electron accelerated through a potential difference of one Volt. $1\text{eV} = 1.602 \times 10^{-19}\text{J}$.

a field theory where the matter fields are required to obey certain symmetry conditions, called ‘gauge symmetries’, then for each symmetry the equivalent of a force will arise. In fact, since these symmetries are imposed locally at each point in space-time, the somewhat uncomfortable Newtonian concept of action-at-a-distance is removed! — a truly elegant solution to the age-old problem of force.

To further illustrate the immense power of symmetry in physics, consider the wealth of electromagnetic phenomena that we know of; that is, *everything* from radio waves to television screens, from computer circuits to the forces that hold atoms together in molecules, from light and sound to the firing of neurons in the human nervous system. The gung ho particle physicist would claim that all this is described by a simple theory, called Quantum Electrodynamics, or “QED” for short. This theory is built on the postulate that there exists a small set of quantum fields (in fact, you need only three, the electron and the up and down quarks) which share a simple kind of gauge symmetry, called $U(1)$. In Quantum Electrodynamics, it is this simple symmetry property of the electron and quark fields which is responsible for the existence of the photon and thereby for all electromagnetic interactions in Nature.

Of course, nobody in their right mind would use the equations of QED to calculate e.g. the structural strains in a skyscraper. Admittedly, the immense difficulty in *solving* the equations when more than a few particles are involved rapidly grows prohibitive. Nonetheless, there is no indication that Nature is not capable of handling her numbers, though it may border on the impossible for us to understand and duplicate her feats. In this work, some such problems will be encountered, and the shortcomings of at least the author will no doubt be plainly visible.

The simplicity of QED aside, electromagnetism will not play a very great role in the papers comprising this thesis. Rather, it is chiefly in the light of problems connected with the weak force that we are prompted to consider supersymmetry in papers 1–4, while some interesting problems connected with the phenomenology of the unruly strong force are studied in papers 2, 5, and 6. However, before we turn our full attention to these aspects, it is worthwhile to begin with a broader consideration of the elementary world view of modern particle physics, to set the stage for the subsequent discussions.



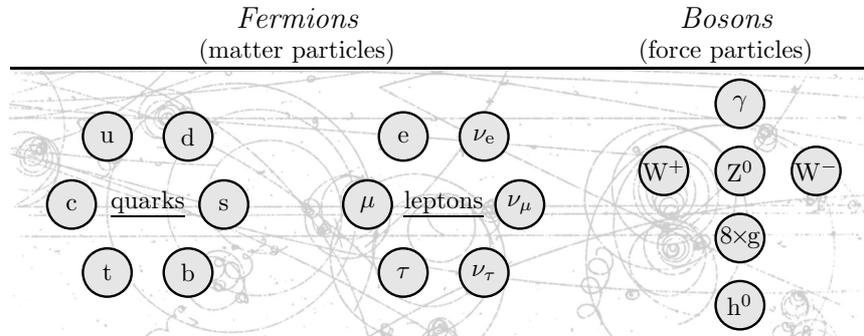


Figure 1: The particles of the Standard Model.

The Standard Zoo

The present state of atomic theory is characterised by the fact that we not only believe the existence of atoms to be proved beyond a doubt, but also we even believe that we have an intimate knowledge of the constituents of the individual atoms...

N. Bohr, Stockholm 1922 [2]

In High Energy Physics, it is a quite unique state of affairs to have a theory of the inhabitants of the fundamental zoo of the Universe survive more than a decade or so. Yet this is precisely the state today. Close to all experiments that have been performed during the last 30 years seem to be in rather perfect agreement with the predictions of a simple theoretical framework built on quantum field theory and using just a small number of fields obeying a small number of symmetries. This enduring model, descriptively known as the Standard Model, or just the SM for short, emerged in the seventies and traces its origins even further back — today it still to most intents comprises our fullest understanding of what goes on in the sub-nuclear world. The inhabitants of the zoo, then as now, encompass the quarks, the leptons, and the Higgs boson, together with the electromagnetic, weak, and strong forces.

An illustration of the particles of the Standard Model is given in fig. 1. Common to all the quarks and leptons is that they are *fermions*. That is, they obey the so-called Pauli exclusion principle, which states that at most one fermion may occupy any given quantum state. This principle is extremely important. For instance, if it were possible for all the electrons of an atom to simultaneously reside in the ground state, there would be no successive filling

of the atomic orbitals and hence no chemistry. The force particles, on the other hand, are examples of *bosons*. The fact that the exclusion principle does not apply to bosons also has farreaching consequences. Without going into details, it is the tendency of bosons to condense into the same quantum state that underlies such spectacular phenomena as superfluidity (liquid flow without resistance) and superconductivity (electric conductivity without resistance).

Returning to the SM, the heaviest known elementary particle is the top quark (t), with a mass almost 200 times that of a hydrogen atom. For comparison, the electron (e) is responsible for only about half a permille of the mass of a hydrogen atom. The neutrinos (ν) were until recently thought to be entirely massless, this also being the prediction of the Standard Model².

Among the force-carrying particles, the photon (γ) and the gluon (g) are massless while the W and Z bosons are about half as massive as the top quark each. The elusive Higgs boson remains a postulate of the theory and has yet to be discovered experimentally. We only know that it should be heavier than about a hundred hydrogen atoms or so, and that our so far splendidly working Standard Model would predict everything to be massless in its absence; hence our candid, if not confirmed, belief in its existence.

Combined, this very modest number of particles and forces successfully describes the vast majority of laboratory experiments to date (excepting of course those dealing with gravity).



So everything is made of quarks and leptons, eh? Who would have *thought* it was so simple?

²It has by now been experimentally established that the neutrinos most likely do have some tiny masses [3]. This must probably be viewed as an indication of physics *beyond* the Standard Model, which we shall return to later.

The Strong Force

So what is the problem? From what has been said so far, the Standard Model appears to be a great theory that works beautifully. The problems of neutrino mass and quantum gravity are easily dealt with, neutrino masses by postulating some invisible particles that lead to no other observable consequences, and quantum gravity by throwing it into a black hole. After all, the product of Newton's G and Planck's \hbar is so extremely tiny that we most likely have many generations to wait before quantum gravity enters the realm of experimental science, if ever.

Nonetheless, particle physics has not entered that decaying stage where only problems of strictly theoretical interest remain to be considered — what was said so far about the simplicity and success of the Standard Model was, fortunately, quite an overglossed picture. There are important and unresolved questions, both within the Standard Model and beyond it, which are investigated actively by experimenters and theorists alike. Among the most notoriously difficult and persevering ones has been the understanding of the dynamics of the strong force.

The strong force, or quantum chromodynamics (QCD) in more modern language, was originally postulated in order to account for the stability of nuclei; since protons tend to electrically repel each other there had to be another, stronger, force at play, that made them stick together in nuclei, which rich imagination dubbed the 'strong force'. However, while its presence is thus revealed to us already at the nuclear level, it is only when probing distance scales smaller even than the proton that the fundamental nature of this force becomes apparent. At such distance scales, the proton has been found no longer to behave like a fundamental particle but rather to be "resolved" into quarks and gluons, which are believed to be the elementary carriers of strong charge.

Interestingly, it is this very duality in the nature of the proton, that it appears as a single entity at low energies while appearing to be resolved into 3 so-called valence quarks at somewhat higher energies, that to the mind of American physicist Murray Gell-Mann was so like the "three quarks for Muster Mark" in James Joyce's *Finnegan's Wake*, that the quarks received their name from there. Concerning the proper pronunciation, Gell-Mann himself does not pronounce "quark" to rhyme with "Mark", whereas Joyce would seem to. In the absence of consensus and hence without any possible justification for further digression, we now return to the strong force.

In analogy with electromagnetism, the strong force is believed to arise as a consequence of symmetry properties of the matter quanta, the quarks. However, in this case the symmetry is not quite as simple as that governing electromagnetism. Instead of just having positive and negative charges, there exist 3 different kinds of strong charge, with corresponding anticharges. These are suggestively labelled *colours* (and anticolours), since they mix with each other in a

way similar to that of ordinary coloured light. Moreover, the force quanta, the gluons, themselves carry strong charge (and even more of it than the quarks) — a fact which causes the strong force to give us considerably larger headaches than simple electromagnetism.

The most immediate consequence may be described as *anti-screening*. Drop a coloured particle in vacuum, and fix it in place. Then watch what happens as you move away from it, carrying an anticharge with you. Anti-screening means that the colour charge will seem to *increase* steadily as you withdraw (more precisely, as the resolution scale is decreased), contrary to what happens for electric charges. In fact, you won't reach many femtometres before the colour field between the two charges narrows into a string-like flux tube of high energy density that pulls at you in an elastic-like manner. At some point, your continued exertions will have pumped enough energy into the field to set the vacuum boiling! Drawing on the colour field energy, virtual quark–antiquark pairs living in the vacuum may make the transition to become ‘on shell’ (i.e. they become as close to real as it has meaning for them to become), in a manner that severs the connection between the two original charges. This mechanism effectively *confines* colour charges to the realm of sub-femtometre distances, since any attempt to separate charges by more than that will immediately be neutralised by the creation of compensating charges from the vacuum. Experimentally, deconfined states have been searched for; an isolated quark, for instance, would stick out like a sore thumb, due to its fractional electric charge. None have yet been found, indicating that confinement is indeed among the stricter of Nature's rules. Interestingly, confinement has never been rigorously proven within QCD, though there are many indications of it.

The second interesting phenomenon is to watch what happens in the opposite case, when you bring two colour charges very close together. The coupling now grows weaker, and the prediction of QCD is that it goes asymptotically to zero as you approach “zero distance” between the two. This property is referred to as *asymptotic freedom*. To aid the inexperienced reader, the behaviour of the strong force field at long and short distances is illustrated schematically in fig. 2 below.

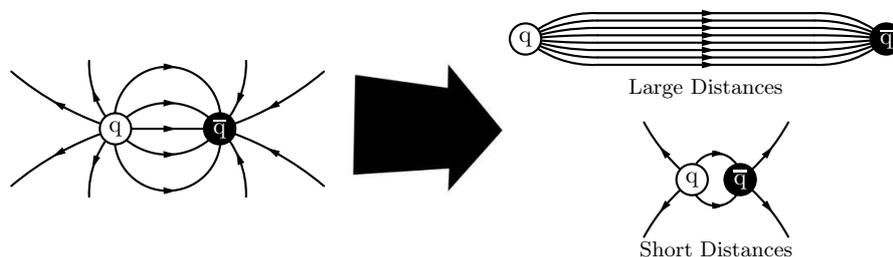


Figure 2: The behaviour of the strong force at large and small distances. The density of the field lines indicates the coupling strength.

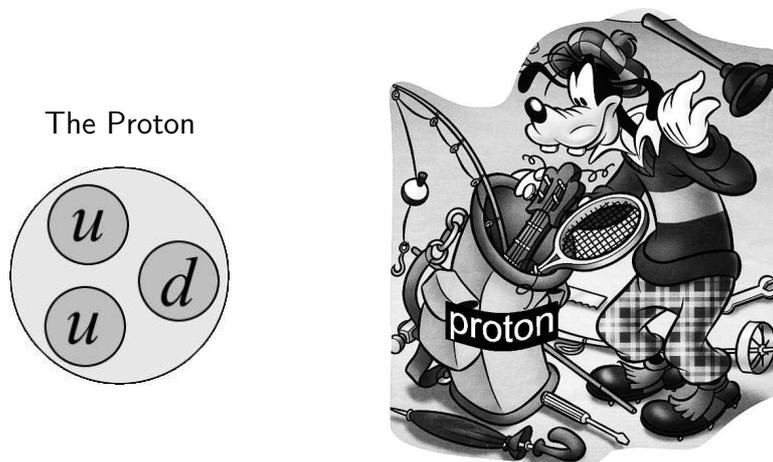


Figure 3: The textbook proton (left) and a more realistic proton (right), the latter shown together with a curious phenomenologist.

Among other things, *asymptotic freedom* guarantees that at high energies (short distances), quarks and gluons behave essentially like free particles, with the strong interaction representing only a small perturbation on that picture. In this case, we can use so-called *perturbation theory* to solve the field equations, an immensely useful and powerful theoretical tool which will be called upon repeatedly in the following. At large distances, however, the *confinement* property mentioned above implies that perturbation theory cannot be used. In this case, we are in fact at a loss for a systematic way of solving the equations. That is, we strongly suspect QCD to be the appropriate theory, but we are unable to solve it for distance scales larger than a Fermi or so (1 Fermi = 1 fm = 10^{-15} m). Nonetheless, there are ways of gaining insight, for instance by the construction of phenomenological models. Such models combine the first principles of the rigorous theory with phenomenological descriptions of the more thorny parts. Comparing different phenomenological models to data and to each other, we gain some insight into the way the theory is working in the problematic regions. The development of certain specific models along these lines is part of the work presented in papers II and V.

Another of the major challenges posed by the strong force is illustrated in fig. 3. To the left is shown the comparatively simple picture of the proton that one might find in an undergraduate textbook, accompanied for instance by the pedagogical statement that the proton consists of three quarks, two up quarks and one down quark, which are held together inside the proton by the strong force. While this picture does carry some truth, a QCD phenomenologist

wishing to determine the proton structure or to consider proton collisions in more detail is faced with a situation more likely resembling that shown to the right of the textbook example. Now, even perturbation theory has its limits, most notably in being tractable only for very small numbers of particles, say a handful or so at a time. A full proton–proton collision, on the other hand, easily involves hundreds of interfering subprocesses at perturbative energies, in addition to the non-perturbative goings-on described above! It is problems of this nature that are studied in papers V and VI.



When I first arrived in Lund, I would come often into Torbjörn’s office, bursting with questions. There would follow the magical incantation “consider e-plus e-minus goes to q q-bar” (in Swedish), whereafter my questions would be answered. (‘q-bar’ here refers to the symbol \bar{q} , used to denote an antiquark.) After a while, I could tell that he thought I had learned something, since he now often began with “consider e-plus e-minus goes to q q-bar gee”, or even “consider p p(-bar) goes to anything”. In the same spirit, the remainder of this brief introduction to the strong force is dedicated to a tour of the various stages and processes that occur in high energy collisions, with the emphasis on those aspects that will be encountered in the papers. However, the task of describing hadronic collisions in all their complexity is a daunting one. The casual reader is therefore encouraged to skip forward to the section on “Beyond the Standard Model”, since the discussion must now become slightly more technical.

*We glibly talk
of nature’s laws
but do things have
a natural cause?*

*Black earth turned into
yellow crocus
is undiluted
hocus-pocus*

Piet Hein, Grooms [11]

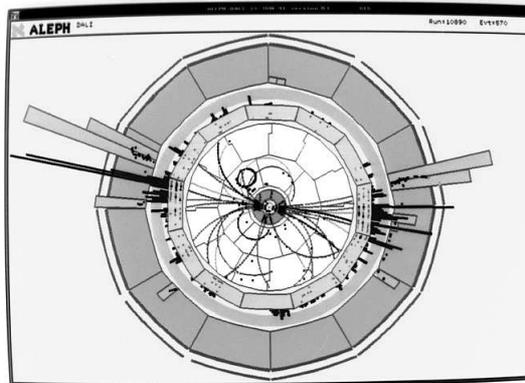


Figure 4: Hadronic decay of a Z^0 in the ALEPH detector. Note that the bending of tracks caused by the strong magnetic field gives an illusion of larger broadening than what is actually the case. This effect is further exaggerated by the fact that the tracks that bend the most are also the ones that carry the least of the total energy.

A High Energy Collision Unravalled

Let us then begin by considering $e^+e^- \rightarrow q\bar{q}$ (in English). A real event of this kind, proceeding via the intermediate creation of a Z^0 boson and recorded by the ALEPH detector at LEP (CERN), is shown in fig. 4. What is observed is far from simply two back-to-back quarks emerging from the interaction point. Rather, two more or less collimated sprays of hadrons may be discerned, called ‘jets’, each carrying an energy equivalent to roughly half the Z mass, $\frac{1}{2}m_Z \simeq 45 \text{ GeV}/c^2$. The jets arise, firstly, due to radiation from the quarks (and, subsequently, from the radiated gluons themselves). The radiation pattern is strongly peaked in the forward direction and tends to produce two jets of well collimated ‘partons’ (quarks and gluons). Secondly, confinement implies that the entire system of colour charges emerging from the bremsstrahlung ‘shower’ must fragment into hadrons. This process adds fine structure and a further amount of broadening to each jet, but since the momentum kicks involved are small, of order a few hundred MeV as compared to $\sim 45 \text{ GeV}$, the overall structure of two well separated jets is maintained.

It should be mentioned, however, that in about 30% of the LEP $e^+e^- \rightarrow Z^0 \rightarrow q\bar{q}$ events a third or even fourth jet is visible. These are due to the occasional hard bremsstrahlung radiation, i.e. a very energetic gluon radiated at large angles off the final state.

Thus, at the very crudest resolution or ‘level of exclusiveness’, there is just the $q\bar{q}$ pair above, and everything else is averaged over. ‘Averaging’ here means that we don’t ask e.g. how many times the quarks radiated or how many pions

were produced. At this level, we measure things like the total cross section as a function of the CM energy, the most inclusive quantity imaginable. At finer resolution scales, we begin to ask increasingly exclusive questions, such as whether there is a third resolvable jet or not. Ultimately, we would like to be able to make predictions down to and even below the hadronisation level.

While it is no problem to calculate rates and distributions for the elementary scattering process $e^+e^- \rightarrow q\bar{q}$ or $e^+e^- \rightarrow q\bar{q}g$ from first principles perturbation theory, the subsequent bremsstrahlung showers and hadronisation phenomena add a level of complexity that is less easily treated, the shower by involving a potentially infinite number of particles (the radiation is not only peaked, it is singular in the forward direction), and hadronisation by being intrinsically non-perturbative.

Concentrating first on the perturbative part, there are essentially three approaches: *fixed order* calculations, *parton shower* models, and *resummation* calculations. For detailed reviews, see e.g. [4]. I will here concentrate on the parton shower approach, which is the one I have been working with and which goes the farthest in obtaining a fully exclusive description of the final state.

As mentioned above, bremsstrahlung processes are strongly peaked in the forward direction. Basically, this is due to the presence of s -channel propagators of essentially massless partons, here light quarks and gluons, but also e.g. electrons in QED, which become singular when the emitted parton is emitted at zero opening angle (*collinear*), and/or when it is low-energetic (*soft*). Thus the soft and collinear limit offers a reasonable starting approximation here. In this limit, the branching processes are furthermore *universal* to a good approximation, i.e. they do not depend on the details of the hard scattering, and they *factorise*, such that an iterative sequence of process-independent branchings may be applied, giving progressively more ‘exclusivity’ to the original hard scattering. The singularities are still there, however, such that an infinite number of branchings would normally result. The reason why parton showers nonetheless remain of practical use is that it is possible to divide the branchings up into a finite number of ‘important’ ones and an infinite number of ‘unimportant’ ones. The key is again *resolution*. It is not possible, for instance, to tell the difference between on one hand a quark and on the other a quark plus an infinite number of gluons (each with infinitesimal energy, such that overall energy and momentum is conserved), unless a minimum of transverse momentum has been imparted to one or more of the gluons. The important point is that the number of branchings that can be resolved above a given scale is always finite. Further, for the case of strongly interacting particles, recall that the basis here is perturbation theory. At resolution scales below ~ 1 GeV or so, perturbation theory breaks down and hadronisation takes over. Thus, the smallest resolution scale, or the largest ‘level of exclusiveness’ that can be reached with the parton shower is roughly the GeV scale.

Turning now to hadronisation, we here rely on phenomenological models

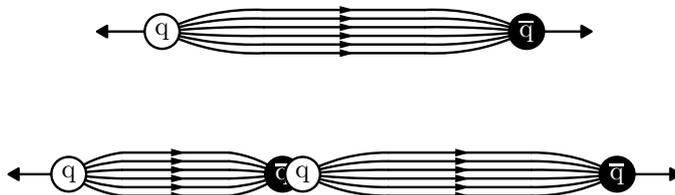


Figure 5: Illustration of a string breakup. A quark and an antiquark emerge from an interaction, tracing a string-like colour flux tube between them. At a separation of a few Fermi, the string breaks, producing a new quark–antiquark pair in the middle which *screen* the original charges from each other, such that there is a region of ‘ordinary vacuum’ in the middle, lowering the total potential energy of the system.

such as the *string* [5] and *cluster* [6] models, or the simpler *independent fragmentation* [7]. Some information can also be gained from *lattice QCD* studies (numerical simulations of QCD dynamics on a space–time lattice). More generally, the topic of hadron dynamics is studied using a variety of approaches, most of which are based on the *operator product expansion*, a systematic if sometimes cumbersome way of constructing effective low–energy theories. Here, we focus on the hadronisation process, leaving hadron decays aside.

Perhaps the most successful hadronisation model has been the Lund string model. This model starts from the assumption of linear confinement, i.e. a linearly rising potential (at long distances) $V(r) = \kappa r$. This form has some support from lattice calculations, and from hadron spectroscopy one finds $\kappa \sim 1$ GeV/fm. The physical picture is that of a narrow flux tube or vortex line being stretched between a colour and an anticolour charge as they move apart. In the Lund model, this is represented by a massless relativistic string of constant energy density κ , spanned between the charges. As the charges separate, the string becomes longer and thus comes to carry more of the total energy. Eventually, there is sufficient energy for the string to break via the production of a new $q\bar{q}$ pair, as illustrated in fig. 5. This process then repeats itself a number of times, until finally a set of string pieces result which all have comparatively small invariant masses (of order hadron masses), such that further breaks will not occur. Such low-mass string systems are in fact nothing but crude representations of mesons, with the quark and antiquark exhibiting a yo-yo like motion about each other inside. Thus, these systems may be mapped onto on-shell hadrons, which may be stable or which may decay further, depending on the type.

Foregoing a longer introduction to the details of the model, it is worth noting its principle features: the fragmentation happens *locally* along a *homogeneous*

string, hence the string breaking probability is to a first approximation constant per unit length per unit time. Nominally, the string breaks are also causally disconnected from each other, i.e. there is no ‘cross talk’ between them. This has important practical consequences, since they can then be considered successively in an iterative scheme. In reality, they do slightly affect each other, for instance by the requirement of overall conservation of momentum and energy, and possibly by collective effects such as Bose-Einstein correlations [8]. Also, the creation of massive quarks in the breakings is suppressed, by the equivalent of a tunnelling factor, i.e. with an exponential dependence on the mass squared. Experimentally, strange quarks are produced less often than u and d quarks, with $P_s/P_u \sim 0.3$, corresponding to a moderate tunnelling suppression. The production of heavier quarks may safely be neglected. The perhaps most striking verification of the existence of a string-like behaviour in Nature was the observation of the so-called *string effect* by the JADE collaboration in the early eighties [9]. The process under study was $e^+e^- \rightarrow q\bar{q}g$, i.e. with an extra gluon in the final state. According to the string philosophy, there should be two hadronising string pieces formed in such an event, one of them spanned between the quark and the gluon, and the other between the antiquark and the gluon. Specifically, there should be *no* string directly between the quark and the antiquark. There should thus be fewer particles produced in the region between the quark and antiquark jets than in the complementary region. Using the fact that the gluon jet most often is the least energetic, JADE was able to statistically separate gluon jets from quark jets, and the resulting measurements confirmed a depletion of hadrons in the expected region, which other models could not account for.

Finally, we now turn to hadron collisions, where the initial state is a conglomerate of strongly interacting partons rather than structureless electrons, and where the final state consists of the hard interaction products plus a number of spectator partons residing in ‘beam remnants’. To begin at the simplest level, the illustration to the right represents a scattering of the elementary quantum states a and b into the final states c and d .



The likelihood for this process to happen, as a function of the a and b centre-of-mass energy $\sqrt{\hat{s}}$, is usually cast in the form of the differential scattering cross section,

$$\frac{d\hat{\sigma}_{ab \rightarrow cd}}{d\hat{t}} = \frac{1}{16\pi\hat{s}^2} |M(ab \rightarrow cd)|^2, \quad (1)$$

where a and b have been assumed massless, $\hat{t} = -Q^2 = (p_a - p_c)^2$ is the four-momentum transfer squared, and $|M|^2$ is the squared sum of amplitudes contributing to the process.

If the incoming particles are *not* fundamental, we may use the *factorisation*

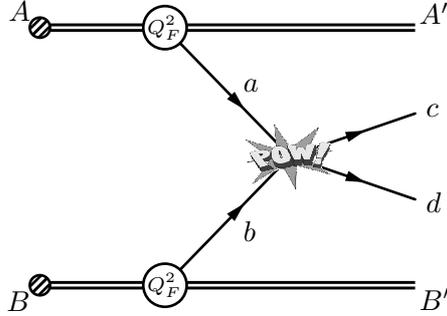


Figure 6: Illustration of eq. (3). The scattering probability is modified to take into account the partonic luminosities inside the parent particles at the factorisation scale Q_F^2 .

theorem. Roughly speaking, this simply amounts to folding eq. (1) with the luminosities of partons a and b inside the incoming particles, call them A and B . More formally, we subdivide the process into a high-virtuality and a low-virtuality part, with the dividing line at the *factorisation scale*, Q_F^2 . The partonic substructures of A and B at that scale, parameterised by so-called parton distribution functions (pdf's), may then be folded into the short-distance cross section to obtain an expression for $AB \rightarrow cd$ (via the subprocess $ab \rightarrow cd$):

$$\frac{d\sigma_{a/A b/B \rightarrow cd}}{d\hat{t}} = \int_0^1 dx_a \int_0^1 dx_b f_{a/A}(x_a, Q_F^2) f_{b/B}(x_b, Q_F^2) \frac{d\hat{\sigma}_{ab \rightarrow cd}}{d\hat{t}}, \quad (2)$$

where $f_{a/A}(x, Q_F^2)$ is standard notation for the pdf giving the probability of resolving a inside A at the scale Q_F^2 , as a function of the momentum fraction $x \sim E_a/E_A$ carried by a . At lowest order in α_s the short-distance cross section $d\hat{\sigma}_{ab \rightarrow cd}/d\hat{t}$ above is given by eq. (1). At higher orders, long-distance contributions must first be subtracted from eq. (1), since by definition all scales below Q_F^2 should be factored into the parton distributions.

The final touch, if we are interested in $AB \rightarrow cd$ via *any* subprocess, is to include a sum over all possible subprocesses (see fig. 6),

$$\frac{d\sigma_{AB \rightarrow cd}}{d\hat{t}} = \sum_{a \in A} \sum_{b \in B} \int_0^1 dx_a \int_0^1 dx_b f_{a/A}(x_a, Q_F^2) f_{b/B}(x_b, Q_F^2) \frac{d\hat{\sigma}_{ab \rightarrow cd}}{d\hat{t}}. \quad (3)$$

This expression is formally all we need to calculate rates and kinematical distributions for c and d production in $A + B$ collisions, provided of course that pdf sets for A and B particles are available.

Notice however, that the factorisation scale Q_F^2 appears explicitly in eq. (3). Rigorously speaking, Q_F^2 is not a physical scale, and so should not appear

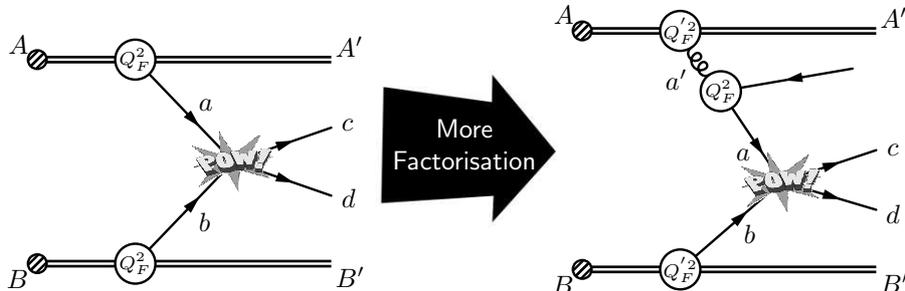


Figure 7: Applying the factorisation theorem at successively smaller factorisation scales yields an iterative scheme for resolving more and more of the quantum structure of the beam hadrons A and B .

in physical observables. Nominally, this scale dependence is cancelled by an implicit scale dependence in the short-distance cross section. At lowest order, this is trivial. With a fixed α_s , the pdf's *scale*, i.e. they depend on x but not on Q_F^2 , and neither is there a scale dependence in eq. (1).

In practice, a certain amount of scale dependence is the price one pays for including higher-order effects incompletely in the calculation. For instance, the use of scale-dependent ‘running couplings’ is standard for obtaining loop-improved tree-level results, since important higher-order contributions are then absorbed into the redefinition of the coupling. While this scale dependence often carries through to the final answer, undesirable from a strictly rigorous standpoint, it is something one can live with since the overall precision of the calculation is still improved.

With a running α_s , the pdf's gain a dependence on Q_F^2 . This scale is then by definition closely tied to the renormalisation scale appearing in $\alpha_s = \alpha_s(Q_R^2)$. Loop calculations indicate that the latter, in its turn, should be connected with the overall resolution scale Q^2 , or better p_\perp^2 [10], of the physics process under consideration. In view of this, the three are often taken to coincide, $Q_F^2 = Q_R^2 = Q^2$ or $Q_F^2 = Q_R^2 = p_\perp^2$.

Finally, by successively using the factorisation property, we may resolve activity at finer and finer resolution scales in the initial state, similarly to how the final state radiation was treated above. The difference is that we must now use parton densities at each step, to describe the changing structure of the beam particle, e.g. the proton, as the resolution scale is made more exclusive. This comes about since we are essentially performing a *backwards* evolution; i.e. starting from a scattering at high virtuality, involving partons selected according to the pdf's at a large factorisation scale, we evolve backwards toward an initial state whose composition is given by the pdf's at a lower scale, as illustrated schematically in fig. 7.

The Monte Carlo Generation

*Whenever you're called on to make up your mind,
and you're hampered by not having any,
the best way to solve the dilemma, you'll find,
is simply by spinning a penny.*

*No - not so that chance shall decide the affair
while you're passively standing there moping;
but the moment the penny is up in the air,
you suddenly know what you're hoping.*

Piet Hein, Grooms [11]

As discussed above, both parton showers and hadronisation models represent essentially process-independent additions to a hard scattering sub-process. As such, both are amenable to an *iterative* formulation, parton showers in terms of successive applications of universal branchings, and hadronisation as a sequence of largely independent hadron formations or string/cluster breakups. The power of computers to perform repetitive tasks according to specific rules readily lends itself to such a formulation.

However, the processes involved are quantum mechanical and hence probabilistic; a given initial state may result in a truly enormous number of different possible final states, depending on the throw of God's proverbial dice. 'Event Generators' essentially simulate the outcomes of such dice throwing, by combining computer generated (pseudo-)random numbers with the statistical Monte Carlo (MC) method (see e.g. [13]). Ultimately, it thereby becomes possible to produce computer simulated 'events'. Fully exclusive descriptions of the final state may thus be achieved, if desired, by combining a simulation of the hard scattering physics with parton shower and hadronisation algorithms (and possibly algorithms describing further physics, such as the occurrence of additional scatterings). An added virtue is that the Monte Carlo method by itself represents the numerical integration technique which converges the fastest for high-dimensional distributions. Thus, an event generator also represents an effective way of calculating cross sections for high-energy processes.

The real boon, however, is that *any* differential distribution of final state observables can be obtained from the generator, since the description is fully exclusive. This is in stark contrast to e.g. resummation descriptions, where one has to proceed observable by observable.

Finally, event generators are used extensively by experimentalists. Since the simulated events are intended to mimic as well as possible those that would be recorded in a real detector, it is possible to pass the events through more

or less sophisticated detector simulations, to gauge the detector response to various physics assumptions, enabling an optimization of the detector design and of the algorithms used to extract physics from the measurements.

Most of the papers involved in this thesis involve the PYTHIA event generator [12], modifying and extending the physics processes that can be described with it in one way or another.

Beyond the Standard Model

There are more things in heaven and earth, Horatio, than are dreamt of in your philosophy.

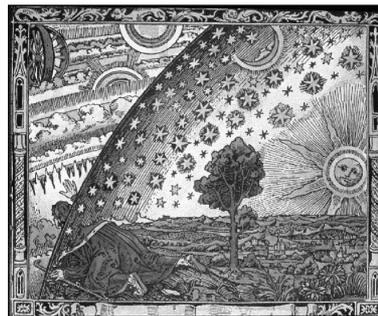
Hamlet (Act I, scene v) [14]

Beyond the Standard Model (BSM); a catch-all phrase used for a broad range of hypothetical ideas, from the already mentioned modest postulate of a few extra neutrinos to theories that have us all living on membranes in higher-dimensional worlds of curled-up extra dimensions.

From the experimental point of view, the search is on for *anything* that might not agree with the predictions of the Standard Model. Scores of measurements are continually being carried out, testing one or another facet of the Standard Model. Meanwhile, within the limits permitted by existing data, theorists speculate.

Despite its many successes the Standard Model is not likely to be ‘the end of the line’; there are missing pieces and internal inconsistencies in the model, which point towards the existence of hitherto undiscovered properties of Nature. Naturally, it is exciting to speculate on what these properties might be. The game is basically to come up with an alternative theory that

1. agrees with present data,
2. ‘solves’ one or more problems in the Standard Model, preferably without introducing new ones,
3. gives predictions which are testable at one or more future experiments,



Looking beyond the Standard Model.

4. is the one for which evidence is found at those experiments, this last condition being the one required to go to Stockholm.

At the most general level, the ambitious and open-minded model-builder has several starting options available. It is possible, for instance, that there could exist new kinds of fundamental matter. Examples here include models with 4 generations of fermions and models with more Higgs bosons. It is also possible that what we *think* of as being fundamental matter, is not. Ideas along these lines are compositeness (substructure in quarks and leptons), technicolour (compositeness of the Higgs boson), and even string theory (where point-like particles are replaced by strings or even membranes).

There could also exist a ‘fifth force’ in Nature, i.e. a new gauge symmetry. Familiar examples are here models with Z' bosons, so-called Left–Right symmetric models, and also technicolour models. The known forces might also turn out all to be but different aspects of one ‘original’ force, as in Grand Unified Theories.

Finally, our concepts of space–time are not inviolate either. There could exist a new kind of space–time symmetry (called supersymmetry), or e.g. relativistic (Lorentz) invariance could be violated to some small extent. Also, when thinking of space and time it is natural to ask why the universe appears to have exactly one time and three space dimensions. The last decade has seen an explosion of activity on this question, with the emergence of many new theories incorporating one or several extra dimensions, which are then curled up in more or less sophisticated ways.

Naturally, some of these ideas are considered ‘more likely’ than others. This is partly due to experimental constraints, item 1 on the list of conditions for success above, and partly due to theoretical considerations. There is also some inevitable prejudice as to which hypothesis is ‘nicest’, aesthetically or otherwise. Giving a satisfactory review of everything that has been proposed and measured is, however, beyond the scope of this introduction. Instead, since papers I–IV concern supersymmetry, I will concentrate on just this possibility in what follows, at what is meant to be a reasonably pedestrian level.

Supersymmetry

The building blocks of matter are fermions, in particular electrons and quarks; the forces, on the other hand, are transmitted via the exchange of bosons, like photons, gluons and the W and Z bosons. Though bosons and fermions to all appearances have very different behaviour, or perhaps precisely because of that, it is tempting to speculate that they might be related at a more fundamental level. Essentially, that is the postulate of supersymmetry (SuSy).

In a supersymmetric theory, each fermion has a boson ‘superpartner’ and vice versa. Each such pair of fields is intimately related by supersymmetry and

is described collectively in so-called *superfields*. If Nature is supersymmetric, the most immediate consequence is thus that there should exist superpartners for all of the Standard Model particles. The naming conventions for these ‘sparticles’ are such that the quarks and leptons have squark and slepton partners, the gauge bosons are accompanied by gauginos, while Higgs bosons give rise to higgsinos. Despite the silly names, these exotic particles currently represent our best guess at detectable New Physics in forthcoming experiments. Conversely, if the searches reveal that such particles really do exist in Nature, and if their properties can be verified to be in accordance with the predictions of SuSy, a completely new type of symmetry of space and time would have been discovered!

While this by itself makes SuSy interesting as a proposition, it was with the discovery that SuSy may solve a fundamental problem known as the *hierarchy problem* that the theory began to gain widespread acceptance. Loosely speaking, every particle gains its mass by coupling to the Higgs field, with the latter supposed even in vacuum to have a constant non-zero strength, a *vacuum expectation value* (vev). We know exactly what the Higgs vev should be, from the spectrum of known particle masses, and even though the Higgs boson itself hasn’t yet been found, we know its mass cannot be much larger. But in the Standard Model, the quantum fluctuations of the Higgs field tend to make the theoretically computed Higgs mass equal to the highest energy scale the theory is postulated to be valid at, regardless of anything else. This is the hierarchy problem in a nutshell; nobody knows what that ‘highest scale’ is for the Standard Model, except that *ultimately* quantum gravity must come into the game, though probably first at scales 10^{17} times larger than the Higgs mass. It would be highly unnatural if the Higgs mass just by accident happened to be 10^{17} times smaller than it had any right to be. Therefore, it is strongly expected that *some* new physics must exist that reduces these large quantum corrections, so that the theoretical prediction becomes more ‘natural’. Tied in with this problem is that it would be nice to have an explanation for the actual *size* of the Higgs mass.

In the beginning of the eighties, a mechanism was discovered, called radiative breaking of electroweak symmetry, whereby SuSy could solve the hierarchy problem and generate the Higgs mass scale in a natural and very elegant way. Add to this that Grand Unified Theories based on the Standard Model were beginning to fail while ones based on SuSy extensions of the Standard Model were seen to work. And add again that many SuSy theories offer an appealing solution to the so-called *Dark Matter* problem in cosmology. This problem has existed since it became clear that the observed amounts of stars and dust are not enough by far to account for how galaxies and galaxy clusters are held together. The universe must contain *a lot* more matter that somehow does not render itself observable to any kind of telescope, hence the name ‘dark matter’. The Standard Model offers no explanation for what the fundamental

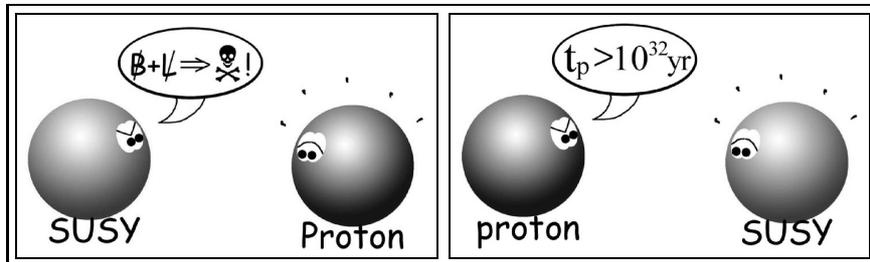


Figure 8: In general, SuSy leads to a rapid rate of proton decay. Experiments, however, give a lower limit on the proton lifetime which is orders of magnitude larger than the age of the Universe. Thus it is SuSy, rather than the proton, that has a problem.

nature of this matter might be, while many SuSy theories predict the existence of particles of the exact type favoured by cosmology, a stable so-called WIMP³. Furthermore, string theory experienced a massive revolution in the mid-eighties, with five theories based on supersymmetric strings emerging as the best candidates for a more all-embracing theory. Combined, all of these indications resulted in an explosive boom of interest in SuSy extensions of the Standard Model. This interest remains today, and searches continue. As the joke goes, there is hope for SuSy, since almost *half* the particles have already been found.

Death to the Proton!

However, SuSy theories do have a few caveats of their own. The first of these concerns the stability of the proton. Since most protons were presumably created during the Big Bang and are still around today, the proton lifetime must be at least of the order of the age of the Universe. In fact, experiments involving extremely large numbers of protons have been able to set limits in excess of 10^{32} years for the longevity of the proton. The Standard Model does not violate this constraint, but only sort of by accident. That is, the stability of the proton is not guaranteed by any fundamental symmetry of the model, but comes about more or less since there just *happens* to be no way the proton could decay via any of the interactions that are possible within the Standard Model.

Unfortunately, the most straightforward SuSy extension of the Standard Model *does* contain such interactions, and without any fundamental symmetry to protect it, the proton would now have a life expectancy of order a fraction

³WIMP: Weakly Interacting Massive Particle

of a second. This situation is illustrated in fig. 8, in which SuSy is informed by the proton that it has a problem.

This problem is not insurmountable, however. One possibility is that proton stability is not quite as accidental as it appears to be. That is, there could exist a more fundamental symmetry which forbids the dangerous interactions. Many examples of such symmetries have been proposed⁴, effectively falling into three categories: conservation of Lepton Number (carried by the charged leptons and by neutrinos), conservation of Baryon Number (carried by quarks), or the simultaneous conservation of both of them.

Normally, SuSy extensions of the Standard Model are based on assuming the conservation of a quantum number called *R*-parity, which is equivalent to assuming *both* Lepton and Baryon number conservation [16]. This gives the simplest phenomenology and has the nice additional property of predicting the existence of the WIMP mentioned above. However, it is possible to construct viable theories where either Baryon Number or Lepton Number is violated, since it is only the *simultaneous* violation of the two that will lead to proton decay. Such *R-parity violating theories* lead to quite different experimental signatures, and also have interesting phenomenological consequences, some of which are investigated in papers I and II.



To round off, this introduction has proven harder to write than I first thought it would be. Nonetheless, it has been interesting and even fun to work on. I hope that reading it may be likewise.

*Under denne Tale var han krøbet i Sengen og slog med begge
Hænder på Dynen, idet han mumlede:
“Haa haa haa! Nu ligger jeg da saa rart og saa godt!”*

Poul Martin Møller, En dansk Students Eventyr [17]

⁴For a discussion of possible proton-protecting symmetries, see e.g. [15].

The Papers

The six articles that constitute the main part of this thesis are here presented in a summary fashion. Each paper is presented individually but with the emphasis on the overall physics picture and on its place in relation to the rest of the articles.

Paper I

When extending the Standard Model with SuSy, the additional constraint of R -parity conservation is usually imposed, this to remove Baryon and Lepton number violating operators from the theory, whose presence would otherwise cause an unacceptably rapid rate of proton decay. However, R -parity is not the only candidate for a proton-protecting symmetry. Viable SuSy theories may incorporate either Baryon or Lepton number violation, since only the simultaneous violation of both leads to proton decay.

In this paper, we explore some phenomenological consequences of Lepton number violation in the Minimal Supersymmetric Standard Model (MSSM). Firstly, we implement rate calculations for a large number of Lepton number violating 2- and 3-body decay channels of sparticles to particles in the PYTHIA event generator. Secondly, some comments are given on experimental triggers and discovery potential at the LHC, using a set of mSUGRA points with various possibilities for the Lepton number violating couplings. To facilitate the latter, a technique based on neural networks is proposed and applied.

Paper II

In R -parity violating SuSy theories, the alternative to Lepton number violation is, obviously, Baryon number violation. From a phenomenological perspective, the latter possibility is more challenging, and more interesting. Where Lepton number is carried exclusively by colour singlets, Baryon number is carried by (fundamental representation) colour triplets, the quarks, and hence colour and colour conservation becomes an important part of Baryon number violating phenomenology.

In this paper, we begin with a natural continuation of the work reported on in paper I, implementing rate calculations for Baryon number violating 2- and 3-body decays in the PYTHIA event generator. The main part of the paper is then dedicated to a discussion of the unique colour topologies involved in these processes, and to the development of a detailed model for non-perturbative physics in their presence. Arguing that, in a colour conserving theory, Baryon number can be viewed as a topological quantum number of the gluon field, we effectively move the Baryon number away from the three quarks and locate it on what we call a ‘string junction’, a topological feature of the gluon field caused by the (antisymmetric) gauge connection between the quarks. From

this standpoint, which allows us to trace the migration of the Baryon number in some detail, we augment the Lund string fragmentation model to include the hadronisation of junction topologies, taking also junction–antijunction systems into consideration.

Lastly, several comparisons are performed with respect to the HERWIG implementation. Mostly, the two codes yield acceptably similar results. The flow of baryon number, however, differs drastically, owing to the lack of a dynamical modelling of this aspect in the HERWIG implementation.

Paper III

In the course of working on the previous papers, the need to translate correctly between the various conventions in use for specifying supersymmetric extensions of the Standard Model had led to recurring, if minor, problems. Specifically, different authors use different parametrisations both of the input assumptions as well as of the resulting mass and coupling spectra. Secondly, SuSy calculations are generally performed in a different regularisation scheme than that used for most other applications. Finally, add to this the more technical aspect of ensuring consistent communication between the large number of SuSy calculational tools presently available, each of which is being continually modified and updated, and one easily realizes that confusion and error must occur, even on the part of specialists.

Thus, aiming to minimise the extent and impact of at least the most trivial sources of error, paper III reports on the development of a standardised set of self-consistent conventions for specifying the most commonly considered SuSy models, combined with a standard set of file structures that SuSy calculational tools may use to interface each other.

Paper IV

In paper I, the possibility of supersymmetry with so-called *bilinear* Lepton number violation was not considered. Nonetheless, this special type of Lepton number violation can have interesting, even spectacular, consequences. Specifically, the lepton number violating couplings and vacuum expectation values induce a mixing between neutralinos and neutrinos, leading to a low-scale see-saw-like mechanism for generating neutrino masses. Indeed, it is possible to construct explicit models that can account for the experimentally observed neutrino mass differences and mixing angles while at the same time leading to testable predictions for the decays of the lightest supersymmetric particle at collider energies.

On this background, paper IV contains a small proof-of-concept study which confirms the possibility of obtaining sufficiently precise measurements at the

LHC to investigate the question of whether neutrino masses could have a supersymmetric origin.

Paper V

The level of sophistication of the junction fragmentation model developed in paper II could perhaps be considered slightly overkill, considering the plausibility of the physics scenario it was applied to, had there not been an ulterior motive in developing it. This motive is furnished by the wish to understand the dynamics of hadron–hadron collisions in more detail, or more precisely, the dynamics of baryon collisions.

Specifically, since hadrons are composite objects, it is possible for several distinct pairs of partons to collide during a single hadron–hadron collision. To the best of our current knowledge, this is not a rare phenomenon. Simple arguments can be used to show that most proton–(anti)proton collisions e.g. at Tevatron or LHC energies are guaranteed to contain several perturbatively calculable interactions. It thus becomes possible for instance that two valence quarks are knocked out of one and the same proton, yielding a situation where the baryon number of the beam cannot naturally be associated with a colour antitriplet diquark, as would otherwise normally be the case. Instead, we must here consider the antisymmetric colour structure of the proton in full, a situation which earlier hadronisation models were simply not capable of addressing.

In paper V we therefore further develop our model for junction fragmentation and apply it to the case of baryon beams (and baryon beam remnants) in the context of multiple parton interactions. The hadronisation description is complemented by several options for the transverse density profile of the incoming hadrons (i.e degree of lumpiness vs. homogeneity) and a sophisticated treatment of correlations in flavour, momentum, and colour between all partons involved, both those undergoing interactions, which are each associated with initial- and final-state parton showers, and those remaining in the beam remnants. Some first studies are carried out, with comparisons to earlier models indicating a good overall description. Some interesting deviations are found, however, indicating that our understanding of in particular the colour correlations is not yet complete.

Paper VI

Consider a hadron–hadron collision at the crudest possible quantum mechanical “resolution scale”. That is, consider an “event” where everything has been averaged over *except* the very hardest parton–parton interaction, say a simple $2 \rightarrow 2$ scattering at high virtuality. Now gradually make the resolution finer. In the ordinary parton shower approach, lower-scale details will now be added to the simple process, an ISR branching here, an FSR branching there, and

finally hadronisation and hadron decays will be resolved as well.

Paper VI, our most recent study, primarily concerns the interplay between initial state radiation (ISR) and multiple interactions (MI). When MI processes are added to the possibilities in the “evolution” described above, it is intuitive that MI and ISR should be competing, in some sense, and that each will affect the other. For instance, if it turns out that a valence quark was knocked out by an interaction at a certain scale, then that should affect the ISR evolution of all other interactions below that scale, simply because that quark is no longer part of the remnant parent hadron at those scales.

We have therefore extended the model proposed in paper V, firstly with p_{\perp} -ordered ISR and FSR algorithms, and secondly with an interleaved evolution (in p_{\perp}) of MI and ISR in parallel. We anticipate that an even more complete model will incorporate the possibility of joining several interactions by perturbative ISR branchings and will return to this in the future.



Acknowledgements

Var vi blevet født paa Venus eller paa Jupiter, vilde vi have maattet leve uden nogensinde at kunne se gennem Skyerne...

Sir James Jeans, *Stjernerne paa Himmelen* [18]

Having always dreamt of becoming an astronomer, and to study the very largest things there are, it is somewhat surprising now to find myself instead studying the very *smallest*. Nonetheless, I am extremely happy that my parents chose to support my interest in the stars rather than try to get me interested in things closer to Earth. For instance, I will never forget when one of the enthusiastic amateur astronomers of KAF showed me the Andromeda galaxy, in my own treasured telescope! Notwithstanding that when I looked in the ocular, I saw nothing but black space on black space, the experience of *almost* seeing something there was *almost* as exciting.

Secondly, my inspiring physics teacher in high school, Niels, had the great idea that more people ought to be exposed to accelerators at a young age, so

the whole class went to see one (no touching!) on a memorable field trip to Aarhus. I left with an impression of complicated electronics, high voltages, and some half-understood nuclear physics, the combined effect of which was quite powerful. At University, I learned classical (Mars!) mechanics from Jens-Martin, the greatest and most inspired teacher I have ever encountered, while Holger in his fashion served as living illustration of quantum mechanics. Once again, I went to Aarhus, this time with Johan and Nino, to Finn Folkmann and his accelerators. I was a teenager let loose in a nuclear physics lab – this time at the controls – and it is simply not possible to describe what that was like. We derived a cross section and got it wrong by a factor 10 billion, but I soon decided I wanted to try my hand at experimental particle physics anyway. I can only hope my precision has gotten better since.

However, in this second resolution I have obviously also failed. This new failure began when Alexander came to Copenhagen. He spent a huge effort trying to beat the theory of B mesons into a couple of inexperienced skulls, one of them belonging to me, the other was attached to Troels Petersen, my office mate. “You have to redo the calculations” became a standard phrase on the corridor, but Alex had an apparently infinite patience. As a CERN summer student, my short career as an experimentalist reached its peak, working on the test beam, Petr and Peter, a few boxes of fast electronics, and an oscilloscope. And there was a girl there who was very sweet, very beautiful, and so smart she could do loop integrals! I fell madly in love with her, and still am.

Well on my way to experimental particle physicisthood, I met Torbjörn and his field, phenomenology, to which also Alex had introduced me: theory that was measurable, testable, and that mattered to experimentalists — it sounded good, and they wanted a new student. And apart from the godawful coffee which I pray I may never taste again, I have had a good time in Lund, with Torbjörn and his family, with Nastyman and the House, the wednesdays at Rydberg’s, the unforgettable performances of THEPISS⁵, and all the beer at Bishop’s. As a young scientist, I am also very grateful to the THEP department for having funded a generous amount of travel, which has allowed me to make so many new contacts and join so many interesting activities abroad, including the highly productive BSM (Bongo-Session-at-Midnight) meeting at Les Houches last year (which in fact laid the foundations for two of the papers in this thesis!). I recall many happy discussions at various blackboards/bars/cliffs with my post doc idols Ben, Peter, Sabine, Werner, David I, David II, and even Tilman. I am also very much indebted to the Royal Physiographic Society in Lund for getting me to Santa Barbara so I could talk to Rick, Joey, and Steve again. There are many more people who deserve thanks, and without whom this thesis would either have been different or my life while writing it more boring. This includes, but is not limited to, Jonas the plant physiologist who proofreads theoretical physics papers at a glance, Paula who organizes the highly enjoyable Nordic

⁵THEPISS: Theoretical High Energy Physics Informal Student Seminars.

LHC Workshops, and the many other people who have supported me or who have invited me to come give a talk somewhere.

First and foremost, I have spent three years becoming a particle physics phenomenologist, and that choice I am sure is among the best and luckiest ones I have ever made.

Wie merkwürdig ist die Situation von uns Erdenkindern! Für einen kurzen Besuch ist jeder da. Er weiss nicht wofür, aber manchmal glaubt er, es zu fühlen. Vom Standpunkt des täglichen Lebens ohne tiefere Reflexion weiss man aber: man ist da für die anderen Menschen —

A. Einstein, *Mein Weltbild* [19]

References

- [1] G. Gamow, *Mr. Tompkins in Paperback* (Cambridge University Press, Cambridge 1993).
G. Gamow, R. Stannard, *The New Mr. Tompkins* (Cambridge University Press, Cambridge 1999).
- [2] N. Bohr, *The structure of the atom*, Nobel Lecture, December 11, 1922.
- [3] Super-Kamiokande Collaboration, Y. Fukuda *et al.*, *Evidence for oscillation of atmospheric neutrinos*, Phys. Rev. Lett. **81** (1998) 1562.
SNO Collaboration, Q. R. Ahmad *et al.*, *Measurement of the charged current interactions produced by B-8 solar neutrinos at the Sudbury Neutrino Observatory*, Phys. Rev. Lett. **87** (2001) 071301.
KamLAND Collaboration, K. Eguchi *et al.*, *First results from KamLAND: Evidence for reactor anti-neutrino disappearance*, Phys. Rev. Lett. **90** (2003) 021802.
- [4] R. K. Ellis, W. J. Stirling, and B. R. Webber, *QCD and Collider Physics* (Cambridge University Press, Cambridge 1996).
CTEQ Collaboration, R. Brock *et al.*, *Handbook of perturbative QCD*, Rev. Mod. Phys. **67** (1995) 157.
- [5] X. Artru and G. Mennessier, *String Model And Multiproduction*, Nucl. Phys. B **70** (1974) 93.
B. Andersson, G. Gustafson and C. Peterson, *A Semiclassical Model For Quark Jet Fragmentation*, Z. Phys. C **1** (1979) 105.

- B. Andersson, G. Gustafson, G. Ingelman, and T. Sjöstrand, *Parton Fragmentation and String Dynamics*, Phys. Rept. **97** (1983) 31.
- B. Andersson, G. Gustafson, and B. Söderberg, *A General Model For Jet Fragmentation*, Z. Phys. C **20** (1983) 317.
- T. Sjöstrand, *Jet Fragmentation Of Nearby Partons*, Nucl. Phys. B **248** (1984) 469.
- B. Andersson, *The Lund Model* (Cambridge University Press, Cambridge 1998).
- [6] R. D. Field and S. Wolfram, *A QCD Model for e^+e^- Annihilation*, Nucl. Phys. B **213** (1983) 65.
- B. R. Webber, *A QCD Model For Jet Fragmentation Including Soft Gluon Interference*, Nucl. Phys. B **238** (1984) 492.
- [7] R. D. Field and R. P. Feynman, *A Parametrization of the Properties of Quark Jets*, Nucl. Phys. B **136** (1978) 1.
- [8] S. Mohanty, *Parton Cascades and Hadronisation in High Energy Processes*, Ph.D. Thesis, Lund University (KFS AB, Lund 2003).
- [9] JADE Collaboration, W. Bartel *et al.*, *Experimental Study of Jets in Electron - Positron Annihilation*, Phys. Lett. B **101** (1981) 129.
- [10] A. Bassetto, M. Ciafaloni, and G. Marchesini, *Jet Structure And Infrared Sensitive Quantities In Perturbative QCD*, Phys. Rept. **100** (1983) 201.
- [11] P. Hein, *Grooks*, in Borgen's Pocketbooks (Narayana Press, Gylling 1995).
- [12] T. Sjöstrand, P. Edén, C. Friberg, L. Lönnblad, G. Miu, S. Mrenna and E. Norrbin, *High-energy-physics event generation with PYTHIA 6.1*, Comput. Phys. Commun. **135** (2001) 238.
- T. Sjöstrand, L. Lönnblad, S. Mrenna and P. Skands, *PYTHIA 6.3 physics and manual*, LU TP 03-38 [hep-ph/0308153].
- [13] F. James, *Monte Carlo Theory and Practice*, Rept. Prog. Phys. **43** (1980) 1145.
- G. Cowan, *Statistical Data Analysis* (Oxford University Press, Oxford 2002).
- [14] W. Shakespeare, *The Complete Works*, ed. W. J. Craig (Magpie Books, London 1993).
- [15] L. E. Ibáñez and G. G. Ross, *Discrete gauge symmetries and the origin of baryon and lepton number conservation in supersymmetric versions of the Standard Model*, Nucl. Phys. B **368** (1992) 3.

-
- [16] P. Fayet, *Mixing Between Gravitational And Weak Interactions Through The Massive Gravitino*, Phys. Lett. B **70** (1977) 461.
G. R. Farrar and P. Fayet, *Phenomenology Of The Production, Decay, And Detection Of New Hadronic States Associated With Supersymmetry*, Phys. Lett. B **76** (1978) 575.
- [17] P. M. Møller, *En dansk Students Eventyr* (Gyldendals Forlag, København 1964).
- [18] J. H. Jeans, *Stjernerne paa Himmelen* (Nordisk Forlag, København 1931).
Original (eng.): *The Stars in their Courses*, Danish translation, F. Topsøe.
- [19] A. Einstein, *Mein Weltbild* (Europa Verlag, West-Berlin 1955).

Searching for L-Violating
Supersymmetry at the LHC

Paper I

Searching for L-Violating Supersymmetry at the LHC

P. Z. Skands

Department of Theoretical Physics, Lund University.
Sölvegatan 14A, S-22362 Lund (zeiler@thep.lu.se)

Abstract The possibility to simulate lepton number violating supersymmetric models has been introduced into the recently updated PYTHIA event generator, now containing 1278 decay channels of SUSY particles into SM particles via lepton number violating interactions. This generator has been used in combination with the ATLFast detector simulation to study the impact of lepton number violation (\cancel{L}) on event topologies in the ATLAS detector, and trigger menus designed for \cancel{L} -SUSY are proposed based on very general considerations. In addition, a rather preliminary analysis is presented on the possibility for ATLAS to observe a signal above the background in several mSUGRA scenarios, using a combination of primitive cuts and neural networks to optimize the discriminating power between signal and background events over regions of parameter space rather than at individual points. It is found that a 5σ discovery is possible roughly for $m_{1/2} < 1$ TeV and $m_0 < 2$ TeV with an integrated luminosity of 30 fb^{-1} , corresponding to one year of data taking with the LHC running at “mid-luminosity”, $\mathcal{L} = 3 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$.

Reprinted from European Physical Journal C, 23, Copyright (2002), with permission from Springer-Verlag.

1 Introduction

Among the primary physics goals of the Large Hadron Collider at CERN, scheduled to turn on some time during 2006, are the exploration of the SM Higgs and top quark sectors, and the direct exploration of the TeV scale with the emphasis on Supersymmetry.

The motivations for believing that Supersymmetry is indeed a property of Nature are many, most importantly the natural and exhaustive extension of the Poincaré group furnished by the Supersymmetric operators [1], the possibility for exact unification of the SM gauge couplings as required by GUTs, the requirement of Supersymmetry for anomaly-cancellation in string theories, and perhaps most importantly, its providing a natural and elegant solution to the problem of scale hierarchy in the SM.

For this latter problem to be solved without unnatural finetuning order by order in perturbation theory, the sparticle masses must lie at or below the TeV scale, and so there is ample reason to believe that SUSY should be observable at the LHC (see e.g. [2]).

The potential for hadron colliders to observe Supersymmetry has been studied in great detail (see e.g. [3, 4]), yet studies tend to concentrate on the simple MSSM framework where both baryon and lepton number conservation are imposed. In section 2 we summarize the theoretical situation, pointing out that it is by no means obvious that baryon and lepton number should be conserved in supersymmetric theories.

The possibility to study lepton and baryon number violating Supersymmetry has already been included [5] in the HERWIG event generator [6] and in SUSYGEN [7], and several studies exist for run II of the Tevatron (for reviews, see [8, 9], latest experimental results [10]) and to a lesser extent for the LHC (see e.g. [11–14]).

However, it has not before been possible to study the full range of decays in these supersymmetric models using the PYTHIA framework (and thus the string fragmentation model), and no dedicated LHC studies have been performed regarding triggers and sensitivity for these theories beyond cases where the LSP was simply forced to decay. In this paper, we go to the general case where all sparticles (excepting the gluino) are allowed to decay via lepton number violating couplings, introducing 1278 lepton number violating decay channels of sparticles to particles in the PYTHIA event generator.

For the purpose of studying the experimental signatures arising from the lepton number violating decay channels, and to quantify to some extent the “discovery potential” of the LHC, some selection of benchmark points is needed. 5 mSUGRA points and, for each of these, 9 scenarios for the 36 \mathbb{Z} coupling strengths are presented and discussed in section 3.

From this point on, the analysis relies on technical aspects of the detector design (or rather, their representation in crude simulation algorithms), and

so we are forced to distinguish between the ATLAS and CMS experiments. Assuming to a first approximation that these experiments will have similar capabilities, we choose to concentrate on ATLAS in the remainder of the paper.

A first step is to define the data sample inside which \cancel{L} -SUSY should be searched for. To this purpose, a selection of triggers designed for \cancel{L} -SUSY scenarios are proposed for “mid-luminosity” running of the LHC ($\mathcal{L} = 3 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$) in section 4.

Before going further, we note that the conventional procedure of studying a few benchmark points in detail is insufficient in the enlarged parameter space opened up by the *L*-violating couplings, even more so since the presence of such couplings also remove important cosmological constraints on the mSUGRA parameters. In particular, the Lightest Supersymmetric Particle (LSP) is no longer required to be neutral [15], and essentially no bound on the mSUGRA parameters can be obtained by requiring a relic density less than or equal to the density of dark matter in the universe.

In section 5, we make a first attempt at coming to grips with the size of the parameter space, proposing a method which relies on neural networks and the grouping of individual scenarios into classes, whereby regions rather than individual points become the studied objects in parameter space. In this work, we define just 2 classes of \cancel{L} -SUSY models, each containing 15 individual scenarios and a class consisting of 5 MSSM scenarios for reference. For each class, a neural network is trained with post-trigger events from all the scenarios in the class against background events, allowing the network to pick out general qualities common to each class without over-fitting to a particular model. Although we study only a few scenarios here, the usefulness of this method should become apparent when considering the requirements posed by more comprehensive scans over parameter space.

After some remarks pertaining to the dangers of using neural networks, we present results for the estimated ATLAS sensitivity, $S/\sqrt{S+B}$, for all scenarios with an assumed integrated luminosity of 30 fb^{-1} . The subsequent step of pinning down the model parameters once a signal has been observed, is clearly more model dependent and is not covered here. An up-to-date review can be found in [9].

A brief outlook and concluding remarks are given in section 6.

2 L-Violating SUSY

In this part of the paper, we do not present new results. Rather, we summarize some theoretical considerations concerning *R*-violation which, although well documented in the literature, may not be in the active memory of all readers.

Despite the many attractive features of Supersymmetry, most notably that it provides a natural solution to the hierarchy problem, it is well known that the most general supersymmetric Lagrangian (containing all terms obeying

$SU(3)_C \times SU(2)_L \times U(1)_Y$ gauge invariance and Supersymmetry) is utterly incompatible with experiment, regardless of whether the assumption of minimal particle content is made or not.

The reason is that baryon and lepton number conservation is not guaranteed by any of the symmetries just mentioned, and the accidental conservation of these quantum numbers in the SM does not hold when extending the SM with Supersymmetry; the full SUSY Lagrangian contains renormalizable lepton and baryon number violating operators [16] which are suppressed only by the SUSY breaking mass scale squared. Unless a much more powerful suppression mechanism is also at work, these operators result in a lifetime for the proton which would be measured in fractions of a second, to be compared with the experimental bound $\tau_p > 1.6 \times 10^{33}$ yr (at 90% CL) obtained by Super-Kamiokande in the $p \rightarrow e^+ \pi^0$ channel [17].

The possibility that the couplings responsible for proton decay are just naturally small but not zero is almost out of the question. Assuming $M_{SUSY} \approx 1$ TeV, the product of the \mathcal{B} and \mathcal{L} couplings involved in proton decay is required to be less than 10^{-25} [18]. Since at least one of the couplings is then forced to have a value below 10^{-12} without any obvious suppression mechanism at work, it is more natural to assume that there is some additional symmetry in the theory, giving zero couplings either for the B -violating terms, the L -violating terms, or both.

Thus, in the Minimal Supersymmetric Standard Model (MSSM), a discrete symmetry, R -parity [19], is customarily imposed which ensures the conservation of baryon as well as lepton number in the supersymmetric Lagrangian. Since there are strong indications that non-gauged symmetries are maximally violated by quantum gravity effects (leading to wormhole-induced proton decay [20]) whereas gauged symmetries are totally stable against such effects [21], we shall here assume that the proton-protecting symmetry is in fact a so-called *discrete gauge symmetry* [22]. In this case, it is of importance to note that although R -parity conservation leads to a comparatively simple phenomenology and a natural dark matter candidate (the LSP), the problem with proton decay is not satisfactorily solved when embedding the supersymmetric theory into more fundamental frameworks (such as GUTs) where baryon and lepton number violation can appear at some higher scale.

In an effective theory valid around the electroweak energy scale, the interactions mediated by super-heavy resonances associated with a higher-scale theory, take the form of non-renormalizable operators (i.e. operators of mass dimension $d \geq 5$) suppressed by $d - 4$ powers of the high scale. Such operators, violating baryon and/or lepton number are generally *not* forbidden by R -parity conservation (see fig. 1), and it has been demonstrated that operators of dimension 5, being suppressed by only one power of the high (e.g. GUT) scale, will cause too rapid proton decay unless their couplings are suppressed by several orders of magnitude [23]. This should, by itself, provide a powerful

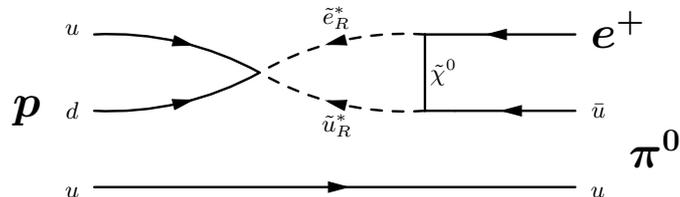


Figure 1: An example of proton decay proceeding via a 5-dimensional operator violating both L and B but conserving R . Due to the Majorana nature of the neutralino, it is drawn without an arrow.

argument for exploring alternatives to R -parity.

In [24], a systematic analysis of discrete gauge symmetries was carried out with the surprising conclusion that only R -parity and a Z_3 symmetry equivalent to L -conservation are anomaly-free, although some model dependence of this analysis [25] implies that also conservation of baryon number could be assured by viable symmetries. We shall here denote the two possibilities by L -parity and B -parity, though this should *not* be taken to mean that these symmetries stand for exact lepton number conservation and exact baryon number conservation (for obvious reasons to do with the matter-antimatter asymmetry of the universe). These names only reflect that in the supersymmetric part of the theory, the corresponding quantities are conserved. They may still be violated by GUT scale operators and will certainly be so by electroweak sphalerons. By far the most interesting observation, however, is that both L -parity and B -parity *but not* R -parity can forbid the dangerous dimension 5 contributions to proton decay.

Since a large number of current GUT and Planck scale theories do contain baryon and lepton number violation exterior to the supersymmetric MSSM framework, which will result in exactly this kind of contributions, it is important that experiments planning to explore SUSY are well prepared for the very distinct experimental signatures which are the hallmarks of R -parity violating scenarios, most importantly the consequences of a decaying LSP, of single sparticle production, and of the removal of the cosmological constraint that the LSP be neutral.

In this article, we assume the framework of the MSSM with the modification that lepton number is violated. This leads to the MSSM interaction Lagrangian

being enlarged by the following terms [18]:

$$\begin{aligned}
\mathcal{L}_{\cancel{L}} = & \lambda_{ijk} (\tilde{\nu}_{Li}^c e_{Lj} \tilde{e}_{Rk}^* + \nu_{Li} \tilde{e}_{Lj} \bar{e}_{Rk} + \tilde{\nu}_{Li} e_{Lj} \bar{e}_{Rk} - i \leftrightarrow j) \\
& + \lambda'_{ijk} \left(\tilde{\nu}_{Li}^c d_{Lj} \tilde{d}_{Rk}^* + \nu_{Li} \tilde{d}_{Lj} \bar{d}_{Rk} + \tilde{\nu}_{Li} d_{Lj} \bar{d}_{Rk} \right. \\
& \quad \left. - \bar{e}_{Ri}^c u_{Lj} \tilde{d}_{Rk}^* - e_{Li} \tilde{u}_{Lj} \bar{d}_{Rk} - \tilde{e}_{Li} u_{Lj} \bar{d}_{Rk} \right) \\
& + \text{h.c.}
\end{aligned} \tag{1}$$

where i, j, k are generation indices (summation implied) and λ_{ijk} is antisymmetric in its first two indices. The terms in the first line of the above equation are customarily denoted LLE terms and the terms in second LQD, in reference to the superfields appearing in the superpotential from which eq. (1) is derived. All of these terms contain two (SM) fermions and one (SUSY) scalar. The oddness under R is thus directly visible here, since the SM fields have $R = +1$ and the SUSY fields $R = -1$. Also, it is clear that each sfermion can decay in a number of ways to two SM fermions via these couplings. L -violating neutralino, chargino, and gluino decays are forced to proceed via one or more intermediate scalar resonances.

A complete list of decay modes and full matrix elements for all decays of MSSM sparticles to SM particles can be found in [26]. Of these, we have implemented all but the gluino decays into the PYTHIA event generator, publically available from version 6.2, with documentation included in the recently revised PYTHIA manual [27]. The conclusions of the present paper are not expected to be changed significantly by the inclusion of \cancel{L} gluino decays since the gluino is typically heavy and thus has a number of other, unsuppressed decay channels available. Details about the PYTHIA implementation can be found in [28].

Note that the sfermion decay matrix elements in [26] have been checked analytically by the author, and that the matrix elements listed there are not directly applicable in PYTHIA. Both PYTHIA and HERWIG (ISAWIG) follow [29, 30] for the chargino and neutralino mixing conventions, but HERWIG uses the opposite convention [31] for the sfermion mixing angles, yielding a relative transposition between the two programs. In addition, extensive counter-checks were made between the two programs to make sure they agree¹. Some non-negligible differences exist:

1. ISAWIG does not use running masses in the evaluation of higgsino-type couplings whereas PYTHIA does. This has been observed to lead to substantial differences between the widths calculated by the two programs (in rare cases factors of 2 or more).
2. α_s at M_Z is used by ISAWIG in the calculation of the gluino decays. In the forthcoming update of PYTHIA where gluino decays are included, α_s

¹This resulted in a few bug-fixes also in the HERWIG code, so ISAWIG has had a few bugs in the ME calculation up to and including v.1104.

will be evaluated at the mass of the gluino.

3. PYTHIA and HERWIG (ISASUSY) do not use the same RGE's, so many parameters can differ quite substantially at the EW scale from the same GUT input, creating an "artificial" dissimilarity between the two programs. Ideally, one should compare models with identical EW scale mass spectra.

3 mSUGRA Models

The models used in this study have not been chosen among the ones initially suggested in the ATLAS Physics TDR [3]. This is partly due to the exclusion of most of these points by LEP (essentially from bounds on the Higgs mass), and partly since it is interesting to enable a direct comparison between the capabilities of the LHC and other, future experiments. The 5 mSUGRA points shown in table 1 have therefore been selected among 14 points which were defined by the CLIC physics study group [32]. More recently - too late to be included in the studies reported here - a new set of standard benchmark points were proposed [2] which have now been adopted by both the CLIC and other linear collider communities. See [2] for mass spectra analogous to table 1. To aid comparison, we briefly list the main differences/similarities between our points and the new benchmark points.

- Point 'A' is a (150GeV) lighter version of P_2 . The total SUSY (pair production) cross section is approximately 4 times larger than for P_2 : $\sigma_A \approx 4\sigma_{P_2}$.
- Points 'B', 'C', 'I', and particularly 'G' are similar to P_9 . $\sigma_{P_9} \approx 0.5\sigma_B \approx 5\sigma_C \approx 2\sigma_I \approx 3.5\sigma_G$.
- Point 'D' is dissimilar to all our points. It is closest to P_2 , but has a 500GeV lighter gluino, and 400GeV lighter squarks which increases the LHC cross section by an order of magnitude.
- Point 'E' is a lower-mass version of F_2 , point 'F' a higher-mass version. $\sigma_{F_2} \approx 13\sigma_F \approx 0.05\sigma_E$.
- Point 'H' is similar to P_7 but slightly heavier, giving a total LHC cross section² $\sigma_H \approx \frac{1}{4}\sigma_{P_7}$.
- Point 'J' is a high- $\tan\beta$ version of P_2 with very close to identical total cross section. The mass hierarchies are also very similar.
- Point 'K' is a large- $m_{1/2}$ sister of the only large- $\tan\beta$ model included in this work, P_{12} . The correspondingly larger masses yield a cross section $\sigma_K \approx 0.04\sigma_{P_{12}}$. The mass spectrum is similar to that of P_7 .

²Point 'H' gave a $\tilde{\tau}_1^+$ LSP when using the GUT input parameters in [2] with ISASUGRA. We therefore chose a 30GeV larger value for m_0 , retrieving $\tilde{\chi}_1^0$ as the LSP.

- We have not included points at $\tan\beta = 45$, but point 'M' can be seen as an even heavier version of point 'K' $\approx P_{12}$ with a total cross section of only 0.1 femtobarns, and point 'L' has a mass spectrum not greatly different from P_9 but with heavier squarks and gluinos giving a cross section an order of magnitude lower than for P_9 .

The total cross sections compared here include all MSSM pair production. Single sparticle production will give additional contributions depending on the strengths of the R -violating couplings.

One should keep in mind that all these points are defined for the MSSM, and as such have neutralino LSP's, a property which it has already been mentioned is not necessary in \mathcal{R} -SUSY scenarios. Though we shall not do so here, it is certainly advisable to explore the experimental consequences of non-neutralino LSP's in more detail.

Both our points and the points in [2] assume a vanishing trilinear coupling at the GUT scale, i.e. $A_0 = 0$. In connection with this work, a small study of the direct consequences of that assumption upon the results presented here was performed. Varying A_0 between 0 and 500GeV for P_2 , P_9 , and P_{12} gave only a weak variation ($\mathcal{O}(5\%)$) of the semi-inclusive L -violating branching ratios (e.g. $BR(\tilde{\chi}_1^0 \rightarrow qq\nu)$), and so apart from the consequences of the distortion of the mass spectrum caused by $A_0 \neq 0$, the main signatures (number of leptons, number of jets, etc.) should be only mildly affected by changes to this parameter.

In addition to the mSUGRA parameters come the 9 λ (LLE) and the 27 λ' (LQD) couplings for which the scenarios listed in table 2 have been studied for each mSUGRA point. The reason we do not consider couplings larger than 10^{-2} is partly due to the present experimental bounds [33] and partly due to the fact that some implicit approximations in this analysis would break down for larger couplings: 1) the \mathcal{R} couplings are not included in the RGE evolution of the SUSY parameters and 2) single sparticle production is not simulated. Note that the large-coupling (a) scenarios only observe the limits on individual couplings given in [33], not the limits on products of the couplings. Thus, these scenarios are useful for studying the consequences of having many large couplings, but one should keep in mind that they are unrealistic in that not *all* of the couplings can be simultaneously large.

On the other hand, if the \mathcal{R} couplings get significantly smaller than 10^{-4} , the LSP lifetime can become so large that it decays outside the detector, mimicking the R -conserving scenarios which have already been extensively studied. For example, for the mSUGRA point F_2 , setting all λ couplings to 10^{-6} and all λ' couplings to zero results in a decay length for the LSP of $\tau c = 40$ m. In the intermediate range, one may see the LSP decay directly inside the fiducial volume of the detector (see e.g. [34]), yet we abstain from relying on such a spectacular signature here so as not to be overly optimistic in our results.

	P₂	P₇	P₉	P₁₂	F₂
GUT Parameters					
$\tan\beta$	5	10	20	35	10
m_0	170	335	100	1000	2100
$m_{\frac{1}{2}}$	780	1300	300	700	600
$\text{sign}(\mu)$	+	+	+	-	+
A_0	0	0	0	0	0
Mass Spectrum					
h^0	118	123	115	120	119
A^0, H^\pm, H^0	1100	1663	416	944	2125
$\tilde{\chi}_1^0$	325	554	118	293	239
$\tilde{\chi}_2^0, \tilde{\chi}_1^+$	604	1025	217	543	331
$\tilde{\chi}_3^0$	947	1416	399	754	348
$\tilde{\chi}_4^0, \tilde{\chi}_2^+$	960	1425	416	767	502
\tilde{g}	1706	2752	707	1592	1442
$\tilde{e}_R, \tilde{\mu}_R$	336	584	156	1031	2108
$\tilde{\tau}_1$	334	574	126	916	2090
$\tilde{e}_L, \tilde{\mu}_L$	546	917	231	1098	2126
$\tilde{\tau}_2$	546	915	240	1051	2118
$\tilde{\nu}$	541	913	217	1095	2125
\tilde{q}_R	1453	2333	612	1612	2328
\tilde{b}_1	1403	2262	566	1412	2010
\tilde{t}_1	1189	1948	471	1241	1592
\tilde{q}_L	1514	2425	633	1663	2343
\tilde{b}_2	1445	2312	615	1482	2310
\tilde{t}_2	1443	2286	648	1451	2018
LHC Parameters					
	P₂	P₇	P₉	P₁₂	F₂
σ_{SUSY} [fb]	130	3.9	24000	110	110
$\times 30\text{fb}^{-1}$	3900	114	720000	3300	3300

Table 1: Selected points of analysis in the mSUGRA parameter space, mass spectra as obtained with ISASUSY, and total SUSY pair production cross section at the LHC.

4 Triggers for \cancel{L} -SUSY

A reasonable aim for the total \cancel{L} -SUSY dedicated trigger rate is about 1Hz. We here focus on rates *after* the events have been filtered through the trigger system, i.e. we make no distinction between trigger levels. This is a technical issue which requires more detailed knowledge of the detector performance at

\cancel{E} COUPLING MODELS			
	a	b	n
		L L E	
λ_{ijk}	10^{-2}	10^{-4}	$\sqrt{\hat{m}_{e_i}\hat{m}_{e_j}\hat{m}_{e_k}}$
λ'_{ijk}	0	0	0
		L Q D	
λ_{ijk}	0	0	0
λ'_{ijk}	10^{-2}	10^{-4}	$\sqrt{\hat{m}_{e_i}\hat{m}_{q_j}\hat{m}_{d_k}}$
		L L E + L Q D	
λ_{ijk}	10^{-2}	10^{-4}	$\sqrt{\hat{m}_{e_i}\hat{m}_{e_j}\hat{m}_{e_k}}$
λ'_{ijk}	10^{-2}	10^{-4}	$\sqrt{\hat{m}_{e_i}\hat{m}_{q_j}\hat{m}_{d_k}}$

Table 2: Selected points of analysis in the $\lambda - \lambda'$ parameter space. The last column corresponds to the “natural coupling” scenario proposed in [23]. $\hat{m} \equiv \frac{m}{v} = \frac{m}{126\text{GeV}}$ and $m_{q_j} \equiv \frac{1}{2}(m_{u_j} + m_{d_j})$. These models (column n) will be referred to as “nLLE”, “nLQD”, and “nLLE + nLQD” in the text below.

mid-luminosity than is currently parametrized in ATLFast. Specifically, no parametrization of the effects of pile-up at mid-luminosity is included, and so we here adopt a best-guess approach, performing the simulation without pile-up, and then multiplying the resulting trigger rates by a factor of 5/3 to estimate the true rates at $\mathcal{L} = 3 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$. This factor is based on the scaling exhibited by the inclusive electron, electron/photon, and $\cancel{E}_T + 2\text{jets}$ trigger rates presented in [3, chp.11] from low to high luminosity.

To retain as much generality as possible in the trigger definitions, it is sensible to use the information contained in the \cancel{E} superpotential terms rather than a selection of decay modes to define the trigger menus. The λ (LLE) couplings are purely leptonic, and thus single sparticle production via these terms is not relevant for the LHC. Thus we are searching for at least two hard leptons (which, however, may be taus), in most cases accompanied by \cancel{E}_T from escaping neutrinos. For the λ' (LQD) terms, we expect at least two leptonic objects accompanied by hard jets when the dominant production mechanism is MSSM pair production. Note that pure jet signatures would also be possible through single slepton production, $q_1\bar{q}_2 \rightarrow \tilde{\ell} \rightarrow q_3\bar{q}_4$. Single squark production via LQD requires a lepton in the initial state and is thus suppressed at the LHC.

The triggers so far investigated are listed in table 3, where background rates and efficiency ranges for all models investigated are shown. For comparison, results for MSSM scenarios are also given. The event generation was performed with an augmented version of PYTHIA 6.155 [35] and ATLFast v.2.53 [36]. With respect to ATLFast, an attempt was made at obtaining more believable muon and electron reconstruction efficiencies by including by hand a muon effi-

Trigger	Background Rate	MSSM Efficiency	LLE Efficiency	LQD Efficiency
mu45I + mu45I	0.2 Hz	1 — 5 %	10 — 40 %	1 — 10 %
e45I + e45I	0.1 Hz	1 — 5 %	1 — 35 %	1 — 10 %
mu15I + e15I	0.1 Hz	2 — 5 %	20 — 60 %	2 — 15 %
mu40I + me75	0.3 Hz	10 — 25 %	40 — 75 %	10 — 35 %
e40I + me75	0.2 Hz	10 — 20 %	15 — 70 %	10 — 35 %
j100 + mu40I	0.5 Hz	10 — 20 %	45 — 70 %	10 — 40 %
j100 + e40I	0.5 Hz	5 — 15 %	15 — 65 %	10 — 35 %
j100 + me175	0.3 Hz	50 — 80 %	35 — 80 %	25 — 80 %
3j50 + mu20I	0.1 Hz	5 — 15 %	45 — 60 %	12 — 40 %
3j50 + e30I	0.1 Hz	5 — 10 %	15 — 55 %	10 — 35 %
3j75 + me125	0.1 Hz	30 — 65 %	30 — 70 %	40 — 90 %
Total Rate	2.1 Hz	60 — 90 %	90 — 99.9 %	60 — 96 %
/Combined Efficiency				

Table 3: Estimated trigger rates for background processes and trigger efficiency ranges for the various MSSM points and \cancel{L} -SUSY scenarios studied for $\mathcal{L} = 3 \times 10^{33} \text{cm}^{-2} \text{s}^{-1}$. The nomenclature follows the ATLAS standard, where e.g. “mu45I” means an isolated muon with $p_T > 45 \text{GeV}$, “me” stands for missing (transverse) energy, and “3j50” means 3 jets, each with $p_T > 50 \text{GeV}$. The total rate is smaller than the sum of the individual rates since there is a certain overlap, and the combined efficiencies can be larger than the individual efficiencies, since there is not a *total* overlap between the triggers.

ciency of 95% independent of muon momentum and an electron reconstruction efficiency of 80% for electrons with $p_T > 50 \text{GeV}$ and 70% for electrons with $p_T < 50 \text{GeV}$. These estimates are based on the ATLAS Physics TDR [3].

Background cross sections, pre-trigger event rates, and sample sizes are listed in table 4. Though low- p_T QCD events have essentially no chance to pass triggers and much less the subsequent analysis cuts, the very high purities required for SUSY signal extraction do not immediately admit these events to be discounted entirely. Rather, a substantial sample of such events was generated with the object of placing an upper bound on the number of low- p_T QCD events remaining after cuts using Poisson statistics. This will be discussed in section 5, however some cautiousness should be employed in interpreting the bounds obtained, since the Monte Carlo is here being stretched far into the tails of its p_T distributions.

We do not consider triple gauge boson production. The cross section (excluding Higgs-induced production) is around 100 fb [37], i.e. at the same level as interesting SUSY cross sections, yet invariant mass cuts on pairs of jets would

EVENT RATES AND SAMPLE SIZES			
Process	σ [mb]	Rate [Hz]	N_{gen}
QCD $2 \rightarrow 2$			
$p_T = 1 - 10\text{GeV}$	55	1.6×10^8	2.5×10^8
$p_T = 10 - 75\text{GeV}$	12	3.7×10^7	2.2×10^8
$p_T = 75 - 150\text{GeV}$	5.5×10^{-3}	1.7×10^4	1.4×10^7
$p_T > 150\text{GeV}$	2.9×10^{-4}	8.7×10^2	1.1×10^7
$t\bar{t}$	6.2×10^{-7}	1.9	5.9×10^6
Z/W	1.2×10^{-3}	3.6×10^3	1.8×10^8
$ZZ/ZW/WW$	1.2×10^{-7}	0.36	5.9×10^6

Table 4: Numbers of generated events for the trigger study. The rates listed are total rates before trigger for $\mathcal{L} = 3 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$. QCD events with the p_T of the hard interaction below 1GeV were not simulated.

presumably be able to reduce this background considerably, and so we do not believe that these processes are dangerous as background sources.

More detailed remarks and a large selection of plots of trigger rates and efficiencies versus thresholds can be found in [28].

From the efficiencies in table 3, one easily sees how much cleaner the signatures of the purely leptonic (LLE) coupling are compared to the signatures involving quarks (LQD) where higher thresholds, due to the hadronic environment, mean smaller efficiencies. From this we conclude that it would be of interest to extend the LQD study, examining whether 2-jet triggers and/or 3-object triggers could enhance the efficiency.

Lastly, though the trigger proposals given here are designed explicitly with \cancel{L} -SUSY in mind, they show a certain overlap with triggers proposed for more conventional physics. The di-muon and 3 jets + lepton triggers, for example, have also been proposed for various Higgs searches. The di-electron trigger as well as 3 jets + electron are proposed to catch $t\bar{t}$ decays. Finally, the conventional SUSY searches also make use of multi-lepton, jets + \cancel{E}_T , and multi-jet signatures [38]. It is therefore not unlikely that the triggers proposed here can be incorporated to some extent into the conventional trigger programme.

5 Discovery of \cancel{L} -SUSY at the LHC

The main purpose here is to deliver an impression of what kind of signal strengths can be achieved at the LHC with 30fb^{-1} integrated luminosity. To maintain generality in view of the more than thousand L -violating channels possible, we do not discuss invariant mass reconstruction or measurements of the

SUSY parameters in general. We focus entirely on the isolation of candidate events through inclusive and kinematic event-shape variables in the attempt to obtain a statistically significant signal as compared with the background expectation after cuts. We do this for several models simultaneously (using neural networks in the final step), an approach which is complementary to the conventional one, where specific scenarios are studied one by one.

Note that no attempt is made, beyond a crude worst-case estimate, to include the effects of pile-up in this analysis.

5.1 Missing Transverse Energy

The post-trigger \cancel{E}_T distribution for the SM and its composition is shown in figure 2a. Note that there are so few double gauge boson events that they are hardly visible on the plot. The peaks at $\cancel{E}_T = 75\text{GeV}$ and at $\cancel{E}_T = 175$ are due to the me75 and me175 triggers becoming active.

In figure 2b the distributions for the most “low-mass” mSUGRA point, P_9 , are shown for $P_9(\text{MSSM})$, $P_{9a}(\text{LLE})$, and $P_{9a}(\text{LQD})$. The degradation of the \cancel{E}_T signature in the \cancel{L} models is evident, though one observes by comparing with fig. 2a that a cut at low \cancel{E}_T values is still possible. Note that the \cancel{E}_T trigger peaks mentioned above are absent for the LLE scenario, this simply because the LLE scenarios do not rely to so great an extent on the \cancel{E}_T triggers, c.f. table 3.

The full range of mSUGRA models are plotted in figure 2c. To enable the models to be shown on the same scale and to be distinguished from each other, the histograms have been normalized to unit area and smoothed. These plots are not intended to give detailed information, but rather to illustrate the spread between the models. Models with heavy squarks and gluinos (P_7 , F_2 , and to some extent P_{12}) have rather flat, MSSM-like signatures whereas the models with lighter sparticles are peaked at low \cancel{E}_T . This is an important point, since the light-sparticle scenarios have high production cross sections and the heavy-sparticle ones low cross sections. A cut on $\cancel{E}_T > 100\text{GeV}$ seems a reasonable compromise between losing events in the peaked distributions (where we have many events anyway) and efficient background rejection required for the heavier scenarios (where we lose little by the cut but have fewer events).

5.2 Leptons and Jets

Due to the fact that the LSP decays, one expects an increased number of leptons and/or jets in the \cancel{L} -SUSY scenarios. In figures 3 and 4 the number of leptons is plotted versus the number of jets for each of the SM backgrounds and for $P_9(\text{MSSM})$, $P_{9a}(\text{LLE})$, and $P_{9a}(\text{LQD})$. All events satisfy $\cancel{E}_T > 100\text{GeV}$.

The absolute normalizations are, of course, very different for each of the SM plots.

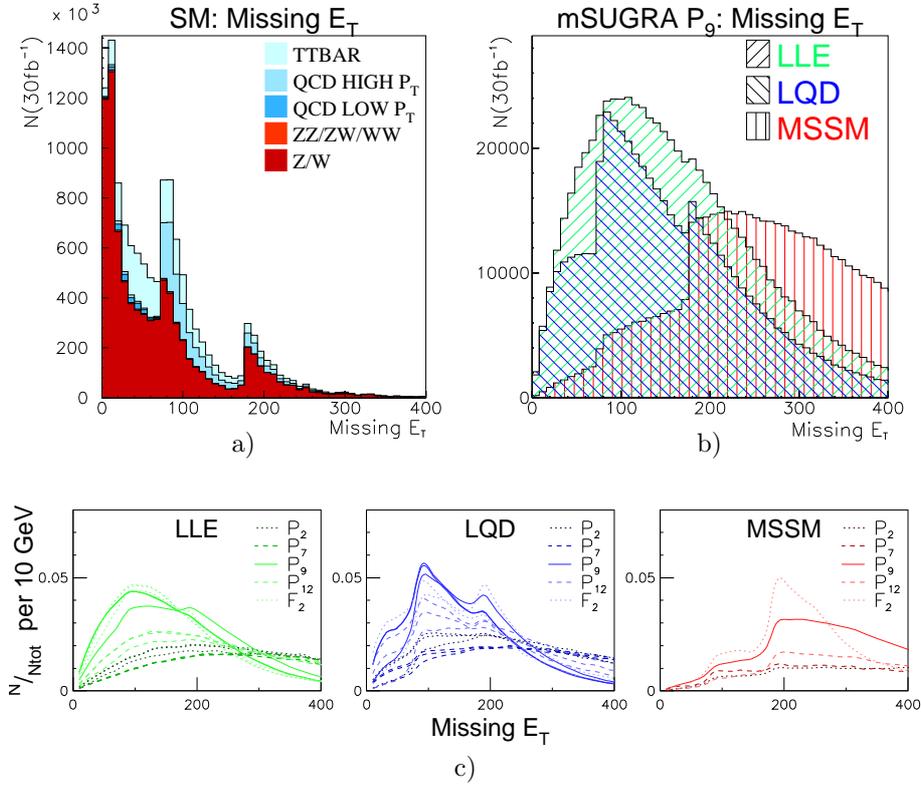


Figure 2: a) and b): \cancel{E}_T signatures for the SM and mSUGRA point 9a models (i.e. all relevant \cancel{R} couplings set to 10^{-2}) normalized to 30fb^{-1} of data taking. “QCD LOW p_T ” means events from the $100\text{GeV} < p_T < 150\text{GeV}$ sample and “QCD HIGH p_T ” events from the $p_T > 150\text{GeV}$ sample. c): Event distributions normalized to unit area for LLE, LQD, and the MSSM for all mSUGRA and \cancel{L} coupling points studied. The last row of plots is not intended to give detailed information, merely to illustrate the spread between the models.

One notices that the number of reconstructed jets goes up to 15 on the plots in figs. 3 and 4. The jet numbers are obtained using the ATLFast standard cone algorithm jet finder, employing a cone size of $\Delta R = \sqrt{(\Delta\phi)^2 + (\Delta\eta)^2} = 0.4$ and with clusters up to pseudorapidities $|\eta| < 5$ included only if their transverse energies are above $E_T^{\text{min}} = 10\text{GeV}$. Whether a resolution of up to 15 jets or more is possible (and reliable) in the finished experiment is questionable, yet the important point remains that there is more hadronic activity associated with the LQD scenarios than is the case for the MSSM, and we should be able to distinguish between the two possibilities using any jet finder.

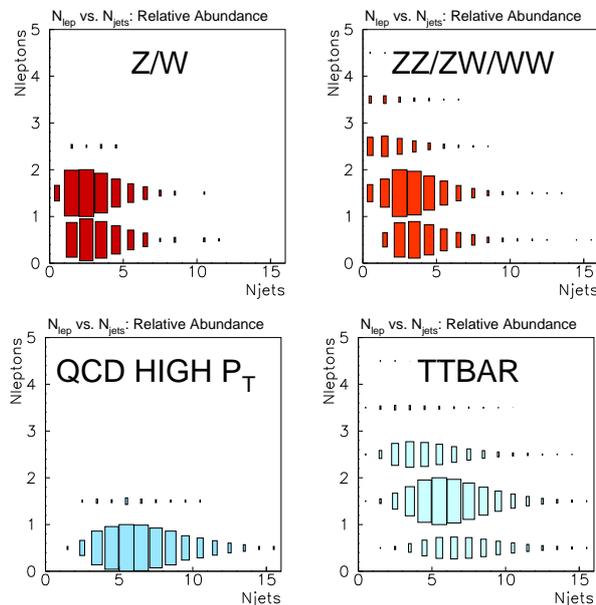


Figure 3: Lepton versus jet multiplicity (see text) for background events surviving the cut on \cancel{E}_T . The numbering of the bins is such that the events with 0 jets are in the bin to the right of the number 0 and events with 0 leptons are in the bin above the number 0.

The increased number of leptons in the LLE scenario as well as the increased number of jets (and a small increase in the number of leptons) in the LQD scenario relative to the MSSM are evident. Since the Z/W and high- p_T QCD backgrounds have the highest cross sections, a cut removing events in the lower left corner was performed, requiring $N_{\text{Jets}} + N_\ell \geq 8$ or $N_\ell \geq 3$. This cut has a rejection factor of about 20 for the SM events and efficiencies above 75% for all \cancel{L} -SUSY scenarios, again with a tradeoff between signal loss at low-mass points and efficient background rejection for high-mass points.

5.3 Additional Variables

In addition to the above mentioned requirements, cuts on several other kinematical variables were performed before the neural network analysis below. We shall here only briefly define and comment these variables (see [28] for details):

- The p_T of the hardest object (jet or lepton) in the event was required to be greater than 200GeV.

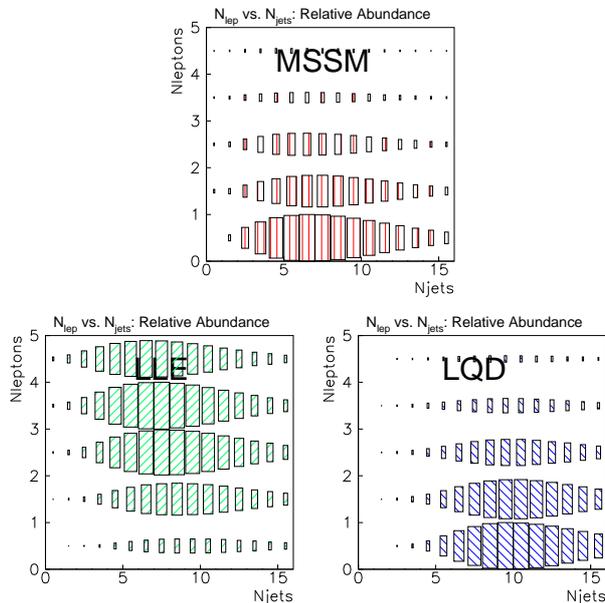


Figure 4: Lepton versus jet multiplicity (see text) in mSUGRA P_9 for events surviving the cut on \cancel{E}_T . The numbering of the bins is such that the events with 0 jets are in the bin to the right of the number 0 and events with 0 leptons are in the bin above the number 0.

- We define the “ p_T -weighted 4-object energy correlation” by:

$$P_{4C} \equiv \frac{1}{12} \left(\frac{E_4}{E_3} + \frac{E_3}{E_2} + \frac{E_2}{E_1} \right) (p_{T1} + p_{T2} + p_{T3} + p_{T4}) \quad (2)$$

where E_1, \dots, E_4 (p_{T1}, \dots, p_{T4}) are the energies (transverse momenta) of the four hardest objects in the event. For pair-produced LSP’s, one would expect four hard objects with more or less equal energies more often than would happen for background events. The energy fractions are weighted by the average p_T to give the variable some extra sensitivity against $t\bar{t}$ events. A cut at $P_{4C} > 200\text{GeV}$ was placed.

- Events with Thrust greater than 0.85 were rejected. The Thrust calculation naturally suffers from the loss of particles down the beam-pipe at high pseudorapidities. Only particles with $|\eta| < 5$ are defined as being inside the active calorimetry range in ATLFAST, and only these particles are included in the Thrust calculation.
- Events with Oblateness greater than 0.4 were rejected³.

³Note that Oblateness is calculated entirely in the transverse plane for hadron colliders.

- Events with Circularity less than 0.1 were rejected⁴.

5.4 Neural Network Cuts

It is clear that the above cuts do not define a dedicated search strategy. They were applied to all scenarios with no distinction between MSSM, LLE, and LQD models. A more comprehensive analysis would of course have to focus on each of these possibilities separately, and also some differentiation between the mSUGRA parameters would be required. In particular sparticles lighter than the top and ones heavier than the top could with advantage be searched for using separate strategies.

This highlights the fact that the mSUGRA parameter space, including now the \mathcal{L} couplings, is not small. It was judged an uneconomical use of resources at this point to optimize a purely physics-based analysis for each of the many possibilities offered by this large parameter space. Instead, three neural networks were trained with post-trigger events to separate MSSM, LLE, and LQD scenarios, respectively, from the background distributions. Each network was thus trained not with one scenario but with equal numbers of events from 15 different scenarios with varying mSUGRA parameters and varying \mathcal{L} -coupling strengths, allowing the networks to pick out general qualities common to each class without over-fitting to a particular model.

This procedure is certainly prone to the dangers which are always present when using neural networks on simulated data, in that the input parameters (*simulated* experimental observables) do not come with warning labels about where in their domains the simulated results can be trusted and where not. However, exercising caution on this point, the method proposed here could be a useful alternative for exploring theories with large parameter spaces.

As inputs were used all of the above mentioned variables along with the p_T of the four hardest jets and the two hardest leptons. The networks employed were single-layer perceptrons using the gradient descent learning algorithm, with biased linear input and output neurons, and hidden neurons activated according to a biased sigmoid ('logistic') function. At the end of each learning period, the networks were trimmed using the OBD (Optimal Brain Damage) prescription [39]. After ten such periods, the networks generally showed a negligible difference in performance on the learning sample as compared with the performance on an independent monitor sample, indicating that over-fitting to the training samples is not a problem.

The training samples consisted of 3000 post-trigger SUSY events, selected equally among the appropriate scenarios, and 4000 background events, weighted according to their post-trigger cross sections to represent approximately 10^7 events in the learning algorithm. An output of zero (one) was the target for background (SUSY) events.

⁴Same comment as for Oblateness.

The finished LLE (LQD) network showed a clear preference for low-Thrust, multi-lepton (multi-jet), low-Oblate-ness events with a mixture of low and high \cancel{E}_T . For slightly higher Thrust values, low lepton momenta were preferred in the LLE case. For the LQD network, the most noteworthy additional feature was a strong relaxation of all other considerations for very high jet energies. Note that the background levels at low Thrust values suffer from some theoretical uncertainty due to the fact that parton showers have been used to generate the multi-jet final states, an approach which in the past has been known to give too few multi-jet events. Though it is not clear that present generators suffer from the same problems, it is important to note that there is a non-negligible and potentially dangerous uncertainty in any analysis relying on parton-shower simulations for large jet numbers and low Thrust values.

Processing now the full event samples remaining after the cuts described above through each of the three (MSSM, LLE, and LQD) networks, a final cut was made requiring outputs larger than 0.9. Since the learning samples were small compared to the full event samples (e.g. about 100,000 SUSY events were generated per scenario and only a few hundred used for training) no effort was made to exclude those events which had participated in the learning process from the analysis. This is, in principle, a source of error, yet we permit ourselves to ignore it due to the smallness of the learning samples and since any problems related to over-fitting have been minimized by OBD. For the SM backgrounds, the generated event samples had been depleted considerably by the initial cuts, and only the $t\bar{t}$ and double gauge boson samples had any events remaining at all after the network cuts (note that the learning samples used for training – for both background and SUSY – were post-trigger events).

The procedure used to construct upper bounds on the actual event numbers was, for N events remaining in a sample, to calculate the mean, μ , of the Poisson distribution which has exactly 5% chance to result in N or fewer events. μ is then interpreted as a “95% CL” upper bound on the number of events which could have passed the cut. For the low- p_T QCD events, this number was then subjected to the same rejection factors under cuts as the high- p_T sample (later the rejections for the $t\bar{t}$ sample were used for both), and the double gauge boson rejection factors were used for the upper bound on single gauge events.

Typically, for 30 fb^{-1} , between 500 and 1000 signal events remain after cuts for P_2 , P_{12} , and F_2 . Many more of course remain for P_9 because of the larger cross section, but there is no hope for P_7 with only 114 events expected in total in 30 fb^{-1} of data. Yet one should bear in mind that single sparticle production which has not been included here could significantly increase the cross section if the \cancel{R} couplings are not much smaller than the gauge couplings. For a hadron machine like the LHC, this effect would only be big for the LQD terms since single slepton resonances would then be possible. Of course, if \cancel{B} terms are present, single squark production would be possible as well.

5.5 Results

We define the statistical significance with which a discovery can be made by

$$P = \frac{S}{\sqrt{S+B}} \quad (3)$$

and conservatively interpret $P > 5$ rather than the conventional $S/\sqrt{B} > 5$ to mean that a 5σ discovery will be possible. Using the event numbers obtained in the analysis (the estimates on B being 95% CL upper bounds) results in the P values listed in table 5.

In reality, P depends on (unknown) systematic QCD uncertainties (parton distributions etc) and should be corrected for the effects of pile-up, and so we can only be confident that a 5σ discovery is possible if P is somewhat larger than 5. Therefore, aside from working with the definition, eq. (3), we attempt to define a more pessimistic quantity in a very crude, *ad hoc* manner which we shall call P_{corr} .

The non-inclusion of pile-up results in too optimistic estimates of S/B . To include an estimate of the reduction of this ratio, we rewrite eq. (3) to:

$$P = \frac{\sqrt{S}}{\sqrt{1+B/S}} \quad (4)$$

where we now include the effects of pile-up by multiplying B/S by some factor. That twice as many background events per signal event could be passing the analysis if pile-up was included seems a reasonably pessimistic guess.

Furthermore, assuming that the intrinsic uncertainty on both S and B coming from uncertainties on QCD parameters will, to a first approximation, work in the same direction and with a comparable magnitude for both B and S , we expect that the denominator in the above formula is not affected by this uncertainty, and so we include the QCD-related uncertainties by reducing the number of signal events in the numerator by a factor of 1.5, believing this to be an adequate worst-case estimate. This yields the following form for the ‘‘corrected discovery potential’’:

$$P_{corr} = \frac{S}{\sqrt{1.5S+3B}} \quad (5)$$

Of course, this quantity should not be taken too seriously. We list it in table 5 merely to show the effects of the stated factors on the discovery potential, i.e. a reduction of S/B by a factor of 2 combined with a reduction of both S and B by $2/3$.

With regard to the sensitivity of this analysis on systematic uncertainties on the normalization of B , note that only for P_9 do we have an $S/B \gg 1$. For the P_2 , P_{12} , and F_2 scenarios S/B lies between one half and unity. We

do not estimate this to be a serious problem since our background estimate is only a 95% CL upper limit (larger event samples would most likely bring the high- p_T QCD component down) and since the background normalization can presumably be determined to better than a factor 2 (as is used in P_{corr}) using complementary regions of phase space where the SUSY population is low or vanishing.

What one can see in table 5 is again that the LLE scenarios typically yield much purer event samples than the LQD ones, due to the hadronic environment at the LHC. It is also worth noticing that it is not easy, from this analysis alone, to discriminate between MSSM, LLE, and LQD scenarios. This is due to the fact that the three networks were trained independently, i.e. they were taught to reject SM events but not events from the other scenario types.

A suggestion for how we might have gained further insight would be to construct *one* network with one output node for each class of models, requiring an output of $(0,0,\dots)$ for background events, $(1,0,0,\dots)$ for events of class 1, $(0,1,0,0,\dots)$ for events of class 2, and so on. One would then be able to construct a measure for the relative probability of the events surviving cuts belonging to a given class of models, by comparing their n -dimensional distribution in the network output space with the expected background distribution.

6 Outlook and Conclusion

6.1 Outlook:

Though some preliminary studies have been performed in the present work, the field is large and many important tasks remain. Recently, a significant theoretical effort has been dedicated to studying the case of baryon number violation [40–42]. Both single sparticle production and baryon number violating SUSY will be included in the PYTHIA generator in the foreseeable future.

This will hopefully spur more LHC activity, with comprehensive studies of both lepton and baryon number violating scenarios. Preliminary studies indicate a lessening of the reach of the LHC in baryon number violating scenarios [40], and it would be of interest to explore this reach with the full range of production and decay mechanisms included. In addition, studies should certainly be carried out for mSUGRA benchmark points without neutralino LSP's.

Finally, triggers dedicated for \cancel{R} -SUSY should be incorporated in the standard ATLAS trigger programme.

6.2 Conclusion:

1278 decay modes of Supersymmetric particles into Standard Model particles through lepton number violating couplings have been implemented in the

ATLAS \cancel{L} -SUSY DISCOVERY POTENTIAL

SUSY Point	NETWORK			SUSY Point	NETWORK		
	MSSM P/P_{corr}	LLE P/P_{corr}	LQD P/P_{corr}		MSSM P/P_{corr}	LLE P/P_{corr}	LQD P/P_{corr}
P_{2a}^{LLE}	24.3/16.2	25.5/17.1	25.4/17.1	$P_{2a}^{LLE+LQD}$	24.3/16.2	25.9/17.4	25.8/17.4
P_{2b}^{LLE}	24.5/16.3	25.8/17.3	25.7/17.3	$P_{2b}^{LLE+LQD}$	24.5/16.3	25.9/17.4	25.9/17.4
P_{2n}^{LLE}	23.2/15.4	24.7/16.5	24.3/16.3	$P_{2n}^{LLE+LQD}$	21.4/14.0	23.2/15.4	24.2/16.2
P_{7a}^{LLE}	0.7/0.4	0.7/0.4	0.8/0.5	$P_{7a}^{LLE+LQD}$	0.8/0.5	0.8/0.5	0.9/0.5
P_{7b}^{LLE}	0.8/0.4	0.8/0.4	0.8/0.5	$P_{7b}^{LLE+LQD}$	0.8/0.5	0.8/0.5	0.9/0.5
P_{7n}^{LLE}	0.7/0.4	0.7/0.4	0.8/0.5	$P_{7n}^{LLE+LQD}$	1.0/0.6	1.0/0.6	1.1/0.6
P_{9a}^{LLE}	191/153	315/256	218/176	$P_{9a}^{LLE+LQD}$	179/143	289/234	203/164
P_{9b}^{LLE}	190/153	316/256	218/176	$P_{9b}^{LLE+LQD}$	178/143	291/236	204/164
P_{9n}^{LLE}	166/133	257/208	169/135	$P_{9n}^{LLE+LQD}$	114/89.3	132/105	111/87.0
P_{12a}^{LLE}	23.4/15.5	25.6/17.2	25.5/17.2	$P_{12a}^{LLE+LQD}$	20.7/13.6	23.8/15.8	24.3/16.3
P_{12b}^{LLE}	23.4/15.5	25.5/17.2	25.5/17.2	$P_{12b}^{LLE+LQD}$	20.8/13.6	23.8/15.8	24.4/16.4
P_{12n}^{LLE}	21.8/14.4	24.2/16.2	24.3/16.3	$P_{12n}^{LLE+LQD}$	16.5/10.6	18.5/12.0	20.8/13.7
F_{2a}^{LLE}	11.3/7.0	14.0/8.8	13.3/8.4	$F_{2a}^{LLE+LQD}$	9.0/5.5	11.9/7.4	11.9/7.5
F_{2b}^{LLE}	11.2/6.9	13.7/8.7	13.1/8.2	$F_{2b}^{LLE+LQD}$	9.0/5.5	11.7/7.3	11.8/7.4
F_{2n}^{LLE}	9.9/6.1	12.4/7.8	12.3/7.7	$F_{2n}^{LLE+LQD}$	7.1/4.3	8.9/5.5	10.4/6.4
P_{2a}^{LQD}	20.9/13.7	24.3/16.2	24.8/16.6	P_2^{MSSM}	10.4/6.4	10.2/6.3	10.2/6.3
P_{2b}^{LQD}	21.4/14.1	24.7/16.5	25.3/17.0	P_7^{MSSM}	0.2/0.1	0.2/0.1	0.2/0.1
P_{2n}^{LQD}	21.5/14.1	23.3/15.5	24.2/16.2	P_9^{MSSM}	136/108	121/95.9	93.5/72.8
P_{7a}^{LQD}	1.0/0.6	1.0/0.6	1.1/0.6	P_{12}^{MSSM}	16.1/10.3	15.5/9.9	16.1/10.3
P_{7b}^{LQD}	1.0/0.6	1.0/0.6	1.1/0.7	F_2^{MSSM}	9.4/5.8	9.7/6.0	10.5/6.5
P_{7n}^{LQD}	1.0/0.6	1.0/0.6	1.1/0.6				
P_{9a}^{LQD}	116/91.6	153/122	125/99.0				
P_{9b}^{LQD}	113/88.7	151/121	123/97.6				
P_{9n}^{LQD}	113/88.5	131/104	110/86.6				
P_{12a}^{LQD}	15.7/10.0	19.2/12.5	20.9/13.7				
P_{12b}^{LQD}	15.8/10.1	19.4/12.6	21.1/13.9				
P_{12n}^{LQD}	16.6/10.6	18.6/12.1	20.9/13.8				
F_{2a}^{LQD}	7.0/4.2	9.5/5.9	10.5/6.5				
F_{2b}^{LQD}	7.0/4.2	9.4/5.8	10.5/6.5				
F_{2n}^{LQD}	6.9/4.2	8.8/5.4	10.3/6.4				

Table 5: ATLAS discovery potential and corrected discovery potential (see text) for all SUSY scenarios investigated using each of the three networks. The numbers shown correspond to an integrated luminosity of 30 fb^{-1} .

PYTHIA event generator. Combining this augmented version of the generator with a crude simulation of the ATLAS detector, trigger menus for mid-luminosity running of the LHC have been proposed and seen to have a high acceptance of supersymmetric events in several L -violating SuperGravity scenarios while still giving event rates in the 1Hz region.

Taking these trigger menus as basis, the possibility for a 5σ discovery after 30 fb^{-1} data taking was estimated for each investigated model, using a technique based on simultaneous classification of groups of models rather than detailed studies of single models.

For cross sections down to 10^{-10} mb (or, roughly, for points with $m_0 < 2\text{ TeV}$ and $m_{1/2} < 1\text{ TeV}$) it was found that a 5σ discovery was possible for all scenarios with 30 fb^{-1} of data. It is estimated that uncertainties related to QCD parameters or pile-up in the detector, both of which have not been taken into account in the present analysis, could not significantly affect this conclusion.

L -violating decays of the gluino and the possibility of invoking B -violation will be implemented in the PYTHIA generator within a foreseeable future.

Acknowledgement: The author gratefully acknowledges support from Frederikke Lørup f. Helms Mindelegat and thanks P. Richardson and S. Mrenna for helpful correspondence.

References

- [1] R. Haag, J. T. Lopuszanski, M. Sohnius, Nucl. Phys. B **88**, 257 (1975).
- [2] M. Battaglia et al. [hep-ph/0106204]
- [3] ATLAS Collaboration, CERN/LHCC 99-14, 1999.
- [4] SUGRA Working Group, S. Abel et al., Report of the sugra working group for run ii of the tevatron, 2000. [hep-ph/0003154]
- [5] P. Richardson, PhD thesis, University of Oxford, 2000. [hep-ph/0101105]
- [6] G. Corcella et al., JHEP **01**, 010 (2001). [hep-ph/0011363]
- [7] N. Ghodbane et al. [hep-ph/9909499]
- [8] S. Ambrosanio et al., in **Tevatron Run II SUSY/Higgs workshop**, 1998. [hep-ph/0006162]
- [9] B. Allanach et al., in **Tevatron Run II SUSY/Higgs workshop**, 1998. [hep-ph/9906224]
- [10] D0, V. M. Abazov et al., (2001). [hep-ex/0111053]
- [11] S. Dimopoulos et al., Phys. Rev. **D41**, 2099 (1990).
- [12] H. Dreiner, P. Richardson, M. H. Seymour, (2000). [hep-ph/0001224]

-
- [13] S. Abdullin et al., in **Physics at TeV Colliders**, Les Houches, 1999. [hep-ph/0005142]
- [14] Y. Xi et al., Phys. Rev. **D64** (2001). [hep-ph/0107006]
- [15] J. Ellis et al., Nucl. Phys. **B238**, 453 (1984).
- [16] S. Weinberg, Phys. Rev. D **26**, 287 (1982).
- [17] Super-Kamiokande Collaboration, M. Shiozawa et al., Phys. Rev. Lett. **81**, 3319 (1998). [hep-ex/9806014]
- [18] H. Dreiner, in **Perspectives on supersymmetry**, ed. G. Kane, chap. 20, (World Scientific, 1998). [hep-ph/9707435]
- [19] G. R. Farrar, P. Fayet, Phys. Lett. B **76**, 575 (1978).
- [20] G. Gilbert, Nucl. Phys. B **328**, 159 (1989).
- [21] L. M. Krauss, F. Wilczek, Phys. Rev. Lett. **62**, 1221 (1989).
- [22] F. Wegner, J. Math. Phys. **12**, 2259 (1971).
- [23] I. Hinchliffe, T. Kaeding, Phys. Rev. D **47**, 279 (1993).
- [24] L. E. Ibáñez, G. G. Ross, Nucl. Phys. B **368**, 3 (1992).
- [25] T. Banks, M. Dine, Phys. Rev. D **45**, 1424 (1992). [hep-th/9109045]
- [26] H. Dreiner, P. Richardson, M. H. Seymour, JHEP **001** (2000). [hep-ph/9912407]
- [27] T. Sjöstrand, L. Lönnblad, S. Mrenna. [hep-ph/0108264]
- [28] P. Skands, Master's thesis, Niels Bohr Inst., 2001. [hep-ph/0108207]
- [29] H. E. Haber, G. L. Kane, Phys. Rep. **117**, 76 (1985), Mixing conventions can be found in App. C.
- [30] J. F. Gunion, H. E. Haber, Nucl. Phys. B **272**, 1 (1986), SLAC-PUB-3404.
- [31] H. Baer, J. Sender, X. Tata, Phys. Rev. D **50**, 4517 (1994). [hep-ph/9404342]
- [32] A. de Roeck, CLIC, Talk given at the CERN EP-TH Faculty Meeting, 2001.
- [33] B. C. Allanach, A. Dedes, H. K. Dreiner, Phys. Rev. D **60**, 075014 (1999). [hep-ph/9906209]
- [34] H. Dreiner, G. G. Ross, Nucl. Phys. **B365**, 597 (1991).
- [35] T. Sjöstrand, Comp. Phys. Comm. **82**, 74 (1994).
- [36] E. Richter-Was et al., ATLAS Internal Note ATL-PHYS-98-131, 1998.
- [37] G. Altarelli, M. Mangano, editors, *Proceedings of the Workshop on Standard Model Physics (and more) at the LHC*, 2000, CERN 2000-004.
- [38] J. Bystricky et al., ATLAS Internal Note (1996), DAQ-NO-054.

- [39] Y. LeCun et al., in **Advances in Neural Information Processing Systems 2**, eds. D. S. Touretsky et al., p. 396, (MIT Press, 1990).
- [40] H. Baer, C. Chen, X. Tata, *Phys. Rev. D* **55**, 1466 (1997). [hep-ph/9608221]
- [41] B. C. Allanach et al., *JHEP* **03**, 048 (2001). [hep-ph/0102173]
- [42] B. C. Allanach et al., *JHEP* **09**, 021 (2001). [hep-ph/0106304]

Baryon Number Violation and String Topologies

Paper II

Baryon Number Violation and String Topologies

T. Sjöstrand¹ and P. Z. Skands²

*Department of Theoretical Physics,
Lund University,
Sölvegatan 14A,
S-223 62 Lund, Sweden*

Abstract In supersymmetric scenarios with broken R -parity, baryon number violating sparticle decays become possible. In order to search for such decays, a good understanding of expected event properties is essential. We here develop a complete framework that allows detailed studies. Special attention is given to the hadronization phase, wherein the baryon number violating vertex is associated with the appearance of a junction in the colour confinement field. This allows us to tell where to look for the extra (anti)baryon directly associated with the baryon number violating decay.

Reprinted from Nuclear Physics B, 659, Copyright (2003), with permission from Elsevier.

¹torbjorn@thep.lu.se

²zeiler@thep.lu.se

1 Introduction

There are good reasons to consider supersymmetry (SUSY) as the next logical step in formulating a consistent theory of particle physics. It complements the Standard Model by the possibility to secure the Higgs potential against excessive radiative corrections, and provides a framework for understanding electroweak symmetry breaking. SUSY is the largest symmetry, within the requirements of conventional field theory, where internal symmetry groups can be combined with the space-time Poincaré group. It thereby points the way towards a unified theory also including gravity, e.g. based on superstrings. For this and other reasons SUSY has been a prime scenario in the search for manifestations of physics beyond the Standard Model. Furthermore, in planning future experiments, SUSY offers a broad range of distinctive signatures that form convenient targets for detector performance and analysis strategies.

In the Minimal Supersymmetric extension of the Standard Model (MSSM), the standard particle content, extended to two Higgs doublets, is doubled up by the presence of superpartners to all normal particles. A multiplicative quantum number called R -parity may be defined by $R = (-1)^{2S+3B+L}$, where S is the particle spin, B its baryon number and L its lepton number, such that normal particles have $R = 1$ while the superpartners have $R = -1$. If R -parity is conserved, supersymmetric particles are pair produced, and the Lightest Supersymmetric Particle (LSP) is stable and hence a Dark Matter candidate. However, there is at present no deep theoretical motivation why R -parity should not be broken. This paves the way both for baryon number violating (BNV = baryon number violation/violating in the following) and lepton number violating (LNV) processes. If both of these are allowed, proton decay would be extremely rapid, unless the relevant couplings are abnormally tiny [1]. As a rule, phenomenology is therefore based on assuming either BNV or LNV, but not both.

The supersymmetric particles that could be produced in high-energy colliders normally are not directly observable: either they decay rapidly or they are weakly interacting and escape detection. With R -parity conserved, the signals would consist of jets, leptons and missing E_{\perp} from escaping neutrinos and LSP's. In scenarios with BNV the main decay product is jets, with only few leptons or neutrinos, and so observability above QCD backgrounds becomes far from trivial at hadron colliders such as the Tevatron or the LHC. In order to carry out realistic studies it is therefore necessary to have a detailed understanding of the properties of both signal and background events. The prime tool for achieving such an understanding is to implement the relevant processes in event generators, where simulated events can be studied with all the analysis methods that could be used on real events. Main generators used for SUSY studies include ISAJET [2], HERWIG [3], SUSYGEN [4], and PYTHIA [5]. Traditionally the PYTHIA framework assumes R -parity conservation, but recently

LNV processes have been implemented and some first studies carried out [6, 7].

In the past, BNV has been modelled [8, 9] and studied [10] in detail in the HERWIG framework, with emphasis on the perturbative aspects of the production process. In this article we present a corresponding implementation in PYTHIA where a special effort is dedicated to the non-perturbative aspects, i.e. what happens with the (anti)baryon number generated by the BNV. This allows us to address the possibility to obtain the “smoking-gun” evidence that a BNV decay has occurred, with questions such as *Could the presence of a violated baryon number be directly observed?* and *If so, what strategy should be used?*. A brief presentation of this work can be found in [11].

In addition, many differences exist between the PYTHIA and HERWIG physics scenarios, for parton showers and underlying events, allowing useful cross-checks to be carried out and uncertainties to be estimated.

The outline of the article is as follows. In section 2, we briefly summarize the short distance physics associated with BNV SUSY and give an account of its implementation into PYTHIA. Next, in section 3, we turn our attention to a parton-shower description of gluon emission in BNV decays. In section 4, the special aspects related to the presence of a net baryon number in the hadronization process are described and an approach based on the Lund model of QCD strings is developed. In section 5 we concentrate on tests of the model, with special attention to the fate of the (anti)baryon produced by BNV. Some (semi)realistic studies are collected in section 6. Finally, section 7 summarizes and gives an outlook.

2 The Baryon Number Violation Scenario

In Supersymmetric theories, it is usually convenient to work at the superpotential level rather than at the Lagrangian level, since the former has a simpler structure and is in any case related to the latter by straightforward manipulations, see e.g. [12]. The most general superpotential which can be written down for the MSSM includes 4 R -parity odd terms:

$$W_{\text{RPV}}^{\text{MSSM}} = \frac{1}{2}\lambda_{ijk}\epsilon^{ab}L_a^iL_b^j\bar{E}^k + \lambda'_{ijk}\epsilon^{ab}L_a^iQ_b^{j\alpha}\bar{D}_\alpha^k + \frac{1}{2}\lambda''_{ijk}\epsilon^{\alpha_1\alpha_2\alpha_3}\bar{U}_{\alpha_1}^i\bar{D}_{\alpha_2}^j\bar{D}_{\alpha_3}^k + \kappa_iL_a^iH_2^a \quad (1)$$

where i, j, k run over generations, a, b are $SU(2)_L$ isospin indices, and $\alpha_{(i)}$ runs over colours. L and Q are the $SU(2)$ doublet (left-handed) chiral superfields for (s)leptons and (s)quarks, respectively, and E , U , and D are the $SU(2)$ singlet (right-handed) superfields for charged (s)leptons, (s)up type, and (s)down type quarks, respectively. The last term represents a mixing between the left-handed (s)leptons and one of the Higgs superfields.

There are two noteworthy aspects about the third term in eq. (1), which we shall henceforth refer to as the ‘‘UDD’’ term: 1) it is the only one which has a non-zero baryon number (all the others violate lepton number), and 2) the colour indices have a Levi-Civita tensor structure. Of course, these two aspects are really one and the same. They express the simple fact that, since SU(3) of colour is unbroken, the only way to make a colour singlet is to combine the three colours antisymmetrically, e.g. as used in constructing ordinary baryon wavefunctions.

In a B -conserving theory like the SM or the R -conserving MSSM, there is no colour antisymmetric perturbative interaction term, i.e. no term with a colour structure like that of the UDD term. Apart from extreme occurrences, like knocking two valence quarks out of the same proton in different directions, by two simultaneous but separate interactions, normal high-energy events would therefore not fully display the antisymmetric colour structure of the proton. Instead, it is normally enough to consider a baryon as consisting of a colour triplet quark and a colour antitriplet diquark, where the internal structure of the latter need not be specified. So what is different about the UDD term is that it allows the production of three colour carriers at large momentum separation, without the creation of corresponding anticolour carriers. It is the necessary SU(3) gauge connection between these three partons that will lead us in the development of the nonperturbative framework.

A further point about the UDD term is that the contraction of the ϵ tensor with $\bar{D}^j \bar{D}^k$ implies that λ''_{ijk} should be chosen antisymmetric in its last two indices, since a (j, k) -symmetric part would cancel out.

The part of the Lagrangian coming from the UDD superpotential term in which we are interested is:

$$\mathcal{L}_{\text{BNV}} = \frac{1}{2} \lambda''_{ijk} \epsilon^{\alpha_1 \alpha_2 \alpha_3} \left(\bar{u}_{R\alpha_1}^i (\tilde{d}^*)_{R\alpha_2}^j (d^c)_{R\alpha_3}^k + \bar{d}_{R\alpha_1}^j (\tilde{u}^*)_{R\alpha_2}^i (d^c)_{R\alpha_3}^k - (j \leftrightarrow k) \right) + h.c. \quad , \quad (2)$$

where we have made the choice of not yet using any of the antisymmetry requirements, so that the ordinary Einstein summation convention applies. Superscript c implies charge conjugation and \tilde{q}^* denotes a charge (=complex) conjugate squark. From this, the possible lowest-order BNV 3-point functions can immediately be read off. In this paper, we only consider sparticle decays; BNV production mechanisms are ignored. This is first and foremost a conservative approach, since any additional sparticle production could only increase the observable signal. Secondly, the underestimation of the sparticle production cross sections is small as long as 1) the BNV couplings are small compared to the gauge couplings and 2) the squarks are not so heavy that single-squark production via BNV is significantly enhanced over the ordinary pair production processes. For discussions of single squark production, see [13].

Combining the vertices in eq. (2) with the full MSSM Lagrangian, also de-

cays involving one or more gauge couplings are clearly possible, e.g. neutralino decay via $\tilde{\chi}^0 \rightarrow \tilde{q}_i (\rightarrow \bar{q}_j \bar{q}_k) \bar{q}_i$. The BNV SUSY decay processes currently implemented in PYTHIA, with Born level matrix elements as calculated by [8], are:

$$1) \quad \tilde{d}_{jn} \rightarrow \bar{u}_i \bar{d}_k \quad (36)$$

$$2) \quad \tilde{u}_{in} \rightarrow \bar{d}_j \bar{d}_k \quad (18)$$

$$3) \quad \tilde{\chi}_n^0 \rightarrow u_i d_j d_k \quad (144)$$

$$4) \quad \tilde{\chi}_n^+ \rightarrow u_i u_j d_k \quad (30)$$

$$5) \quad \tilde{\chi}_n^+ \rightarrow \bar{d}_i \bar{d}_j \bar{d}_k \quad (14)$$

$$6) \quad \tilde{g} \rightarrow u_i d_j d_k \quad (36)$$

where n runs over the relevant mass eigenstates: $n \in \{L, R\}$ for the first two generations of squarks, $n \in \{1, 2\}$ for the third generation squarks and the charginos, and $n \in \{1, \dots, 4\}$ for the neutralinos. The numbers in brackets are the number of modes when summed over n , i , j , and k , and over charge conjugate modes for the Majorana particles.

The matrix elements for these processes, as implemented in PYTHIA, are not quite identical to those used in HERWIG. Most importantly, PYTHIA uses running masses and couplings at some points where HERWIG does not. See [7] for a list of these differences.

When calculating the partial widths (and hence also the rates) into these channels, we integrate the matrix elements over the full phase space with massive b and t quarks, and massive sparticles. All other particles are only treated as massive when checking whether the decay is kinematically allowed or not, i.e. they are massless in the phase space integration.

A feature common to both programs is how double-counting in the BNV three-body modes is avoided. The diagrams for these modes contain intermediate squarks which may be either on or off the mass shell, depending on the other masses involved in the process. If a resonance can be on shell, we risk doing double counting since PYTHIA is then already allowing the process, in the guise of two sequential $1 \rightarrow 2$ splittings. Technically, this means that the list of $1 \rightarrow 3$ BNV widths obtained by a call to PYSTAT(2) only represent the non-resonant contributions, the resonant ones being accounted for by sequences of $1 \rightarrow 2$ splittings in other parts of the code.

In the description of the momentum distribution in a three-body resonance decay, the default PYTHIA procedure is to assume an isotropic phase space, i.e. the matrix-element information used above to obtain partial widths is neglected here. This should not be a bad approximation when the intermediate squark propagators are far from mass shell over the full phase space and therefore do not vary much, but could be a problem when a squark mass (plus the mass of the corresponding quark) is only slightly above the gaugino one. In section 5 we return to this issue, comparing with HERWIG, where the full, unintegrated matrix elements are used to give a more correct phase space population.

3 Parton Showers

The production and decay of unstable particles is normally associated with bremsstrahlung emission of gluons and/or photons as applicable. This radiation is conveniently described in the parton-shower language, where the effects of multiple emissions are explicitly included by iterative applications of the relevant splitting kernels, such as $q \rightarrow qg$, $g \rightarrow gg$, $g \rightarrow q\bar{q}$ and $q \rightarrow q\gamma$. Even though existing showering algorithms typically only fully resum the leading logarithms, most of the effects of next-to-leading logs are also included in the form of energy-momentum conservation, recoil effects, scale choice in α_s , coherence effects, and matching of energetic emissions to higher-order matrix elements. At LEP the PYTHIA, HERWIG and ARIADNE [14] radiation routines have been well tested. All three are found to describe the data well, although some problems exist [15]. Since ARIADNE has not been used for SUSY studies, we restrict the continued discussion to the former two.

The HERWIG algorithm is based on an evolution variable that ensures angular ordering; subsequent emissions occur at smaller and smaller angles. Thereby coherence effects are respected, i.e. double-counting of emission is avoided. The colour flow of the hard process defines cones around the partons, within which all emissions occur. In the limit where each emitted gluon is much softer than the hard-scattering partons, and much softer than all preceding gluons emitted, this approach can be shown to give the correct emission rate, and also when going away from this limit it should provide a good overall description. However, a consequence of the way the kinematics of the algorithm is constructed is that “dead zones” occur, within which no emission at all is possible [16]. Thus, in $e^+e^- \rightarrow q\bar{q}g$ events, the region of energetic gluon emission at large angles to both the q and \bar{q} is completely empty. The solution [16] is to combine two classes of events, both $q\bar{q}$ and $q\bar{q}g$ ones, with the latter picked in the “dead zone” region of the three-body phase space according to the relevant matrix elements. All events, i.e. from either class, are then allowed to shower further. However, such corrections have only been worked out for a few cases, e.g. $e^+e^- \rightarrow q\bar{q}g$ and top decay [17].

The PYTHIA algorithm is based on evolution in the virtuality of the emitting parton, with the requirement of angular ordering enforced by vetoing non-angular-ordered emissions. This gives a less exact description of coherence in the soft-gluon limit. Yet it does allow the full three-body phase space to be populated, simply by not imposing any angular ordering constraint for the first emission of each of the two original partons. As it turns out, the incoherent addition of radiation from these two sources tends to give an overestimation of the $q\bar{q}g$ rate in the hard, non-collinear gluon region, while the soft and collinear regions give what they should. It is therefore straightforward to introduce a correction factor, whereby a fraction of the emissions are vetoed, so that the remaining fraction agrees with the desired three-jet rate [18]. Some years ago,

the relevant gluon-emission corrections were calculated and implemented for most two-body decays within the SM and the MSSM, with R -parity conserved [19].

Normal processes are characterized by unbroken colour lines: the colours present in the initial state will also be there in the final state, with the proviso that opposite colours and anticolours may be pair-produced or annihilated. This information is used in parton showers, e.g. to set up emission cones. BNV processes are different in this respect: in a $\tilde{q} \rightarrow \bar{q}q$ branching a blue colour line may end and antired and antigreen begin. Therefore the standard rules need to be supplemented. In HERWIG an effort has been made to study the different new topologies and use the soft-gluon eikonal, i.e. spin-independent, expressions to set up the relevant maximum emission cones in the different processes [8]. Still, the “dead zone” issue is not addressed, meaning that the rate of energetic, wide-angle gluon emission is almost certainly underestimated.

By contrast, we shall here allow emission over the full phase space, like before. However, not having calculated the corresponding matrix-element correction factors, one should not expect the correct rate in the hard-gluon region. Technically, in the lack of further knowledge, the $\tilde{q} \rightarrow \bar{q}q$ process is corrected to the eikonal answer for a colour singlet decaying to a triplet plus an antitriplet [19]. For three-body decays such as $\tilde{\chi} \rightarrow qq\bar{q}$ or $\tilde{g} \rightarrow qq\bar{q}$, no matrix element corrections are available.

With three (or more) primary partons to shower, one is left with the issue how the kinematics from the on-shell matrix elements should be reinterpreted for an off-shell multi-parton configuration. We have made the arbitrary choice of preserving the direction of motion of each parton in the rest frame of the system, which means that all three-momenta are scaled down by the same amount, and that some particles gain energy at the expense of others. Mass multiplets outside the allowed phase space are rejected and the evolution continued.

In principle, radiation in the initial and final states may interfere, thereby requiring the inclusion of further coherence conditions. The extent of such interference critically depends on the width of the intermediate resonance: only gluons with energies below this width are emitted coherently from initial and final partons [20]. We here assume that the resonances are sufficiently narrow that such interference effects can be neglected; see further the discussion in section 4.5.

The bottom line is that PYTHIA likely overestimates the hard-gluon emission rate, whereas HERWIG underestimates it. This makes it interesting to compare the two descriptions, and possibly to use differences as a first estimation of systematic uncertainties in our current description of BNV processes.

4 String Topologies

Up till now we have considered short-distance processes, where perturbation theory provides a valid description in terms of quarks, gluons and other fundamental particles. At longer distances, the running of the strong coupling α_s leads to confinement and a breakdown of the perturbative description of QCD processes. Instead there is a transition to a nonperturbative régime, characterized by hadrons rather than by partons. In the lack of an exact approach, this hadronization process must be modelled. The perhaps most successful and frequently used approach is the Lund string fragmentation model [21].

This approach has not before been applied to the colour topologies encountered in BNV. We therefore here extend the string model by the introduction of a junction, where three string pieces come together. Effectively, it is this junction that carries the (anti)baryon number that is generated by a BNV process. The hadronization in the region around the junction will therefore be of special interest.

4.1 The normal string

In this subsection we summarize some properties of the ordinary Lund string model, highlighting the concepts that will be needed for the developments in the remainder of the article. Readers who are already familiar with the fundamentals of the string model may wish to proceed directly to the next subsection.

To illustrate the string model, it is useful to start with the simplest possible system, a colour-singlet $q\bar{q}$ 2-jet event, as produced in e^+e^- annihilation. Here lattice QCD studies lend support to a linear confinement picture in the absence of dynamical quarks, i.e. in the quenched approximation [22]. Thus the energy stored in the colour dipole field between a charge and an 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 non-Abelian character of QCD. The dynamical mechanisms involved are not fully understood, however.

The assumption of linear confinement provides the starting point for the string model. As the q and \bar{q} partons move apart from their common production vertex, the physical picture is that of a colour vortex line, or maybe a colour flux tube, being stretched between the q and the \bar{q} . (The difference between these two terminologies is related to whether the QCD vacuum more resembles a type II or a type I superconductor. This is an interesting question by itself, with the vortex line analogy somewhat favoured by theoretical prejudice, but it will not be crucial for the continued discussion here.) 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 approach 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 parameterizing the position of the axis of a cylindrically symmetric flux tube or vortex line. From hadron spectroscopy, the string constant, i.e. the amount of energy per unit length, is deduced to be $\kappa \approx 1$ GeV/fm.

For a massless $q\bar{q}$ pair moving out along the $\pm z$ axis, the original energy and momentum of the quark is then reduced with time according to

$$E_q(t) = p_{zq}(t) = E_q(0) - \kappa t, \quad (3)$$

with a corresponding equation for $E_{\bar{q}}(t) = -p_{z\bar{q}}(t)$. This classical equation of motion obviously does not take into account the quantum mechanical uncertainty relation; in this sense one may conceive of the (unmeasured) space-time picture as a subordinated intermediate step in the derivation of the (measured) energy-momentum picture.

If the quark instead has a nonvanishing rest mass m_q the equation of motion is

$$\frac{dp_{zq}(t)}{dt} = \frac{dp_{zq}(t)}{dE_q(t)} \frac{dE_q(t)}{dt} = \frac{E_q(t)}{p_{zq}(t)} \left(-\kappa \frac{dz}{dt} \right) = \frac{E_q(t)}{p_{zq}(t)} \left(-\kappa \frac{p_{zq}(t)}{E_q(t)} \right) = -\kappa, \quad (4)$$

which has the solution

$$p_{zq}(t) = p_{zq}(0) - \kappa t; \quad (5)$$

$$E_q(t) = \sqrt{E_q^2(0) - 2\kappa t p_{zq}(0) + \kappa^2 t^2}. \quad (6)$$

A key aspect, which we will make use of for the junction phenomenology below, is that the momentum loss per unit of time is independent of the quark mass.

In the back-to-back configuration the string does not carry any net momentum, but acquires energy proportional to its total length as the q and \bar{q} move apart. Once the endpoint quarks have lost their momenta, they will be pulled back by the string tension and re-acquire momenta in the opposite direction, at the same time as the string length shrinks back. When the quarks meet again the force flips sign and a second half-period of the oscillation begins.

Such stable “yo-yo” modes are used as simple representations of mesons. For a high-invariant-mass $q\bar{q}$ system, the fragmentation into lower-mass mesons proceeds through the production of new $q'\bar{q}'$ pairs, corresponding to dynamical quarks in lattice QCD. Thereby the original system first splits into two colour-singlet systems, $q\bar{q}'$ and $q'\bar{q}$. If the invariant mass of either of these string pieces is large enough, further breaks occur. In the Lund model, the string break-up process is assumed to proceed until only on-mass-shell hadrons remain, each

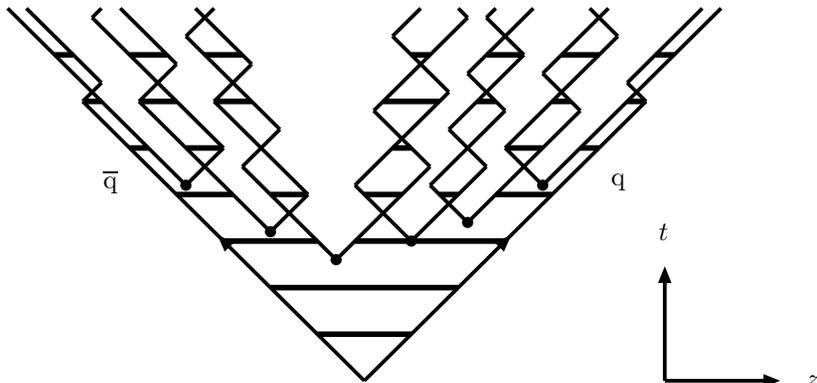


Figure 1: The breakup of an original $q\bar{q}$ system into a set of mesons, each represented by a yo-yo state. For simplicity quarks are here assumed massless, and so move along lightcones. The broken horizontal lines illustrate the string pieces at a few discrete times.

hadron corresponding to a small piece of string with a quark in one end and an antiquark in the other. This process is illustrated in Fig. 1.

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. One therefore is led to formulate an iterative scheme for the fragmentation, as follows. Assume, as above, an initial quark q moving out along the $+z$ axis, with the antiquark going out in the opposite direction. By the production of a $q_1\bar{q}_1$ pair, a meson with flavour content $q\bar{q}_1$ is produced, leaving behind an unpaired quark q_1 . A second pair $q_2\bar{q}_2$ may now be produced, to give a new meson with flavours $q_1\bar{q}_2$, etc. At each step the produced hadron takes some fraction of the available energy and momentum. This process may be iterated until all energy is used up, with some modifications close to the \bar{q} end of the string in order to make total energy and momentum come out right, since all hadrons are required to be on mass shell.

The choice of starting the fragmentation from the quark end is arbitrary, however. A fragmentation process described in terms of starting at the \bar{q} end of the system and fragmenting towards the q end should be equivalent. This “left–right” symmetry constrains the allowed shape of the fragmentation function $f(z)$, where z is the fraction of the remaining light-cone momentum $E \pm p_z$ (+ for the q jet, – for the \bar{q} one) taken by each new particle. The resulting “Lund symmetric fragmentation function” has two main free parameters, which are determined from data.

Viewed in time, the fragmentation process actually starts near the middle of the event and spreads outwards. The $q'\bar{q}'$ production vertices on the average occur along a hyperbola of constant invariant time τ , $\tau^2 = t^2 - z^2$, so even if the string is boosted along the z axis it is still the slow particles in this new frame that are the ones produced first. In this sense, a Lorentz-covariant “inside-out” cascade can technically be described by an “outside-in” iteration scheme. As an order-of-magnitude, $\langle\tau\rangle \approx 1.5$ fm or $\langle\kappa\tau\rangle \approx 1.5$ GeV.

In order to generate the quark–antiquark pairs $q'\bar{q}'$ which lead to string break-ups, the Lund model invokes the idea of quantum mechanical tunnelling. This gives a flavour-independent Gaussian spectrum for the p_\perp of $q'\bar{q}'$ pairs. Since the string is assumed to have no transverse excitations, this p_\perp is locally compensated between the quark and the antiquark of the pair. The total p_\perp of a hadron is made up out of the p_\perp contributions from the quark and antiquark that together form the hadron. Some contribution of soft unresolved perturbative gluon emission may also effectively be included in this description.

The tunnelling picture 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, but only in perturbative parton-shower branchings $g \rightarrow q\bar{q}$.

When the quark and antiquark from two adjacent string breaks are combined to form a meson, it is necessary to invoke an algorithm to choose between the different allowed possibilities, notably between pseudoscalar and vector mesons. Here the string model is not particularly predictive. Qualitatively one expects a 1 : 3 ratio, from counting the number of spin states, multiplied by some wave-function normalization factor, which should disfavour heavier states.

A tunnelling mechanism can also be used to explain the production of baryons. 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 [23]. A more complex scenario is the “popcorn” one [24], where diquarks as such do not exist, but rather quark–antiquark pairs are produced one after the other. This latter picture gives a less strong correlation in flavour and momentum space between the baryon and the antibaryon of a pair.

In all fairness, it should be said that the description of baryon production is one of the less predictive aspects of the Lund model, with many free parameters for the flavour composition that need to be determined from the data. Most single-baryon spectra are well described after such a tuning, except possibly at very large momentum fractions [25]. A reasonable description is also obtained for baryon–antibaryon correlations, although with some disagreements. For the aspects that we will study here, and within the precision allowed by other considerations, the baryon production model should be fully adequate, however.

If several partons are moving apart from a common origin, the details of

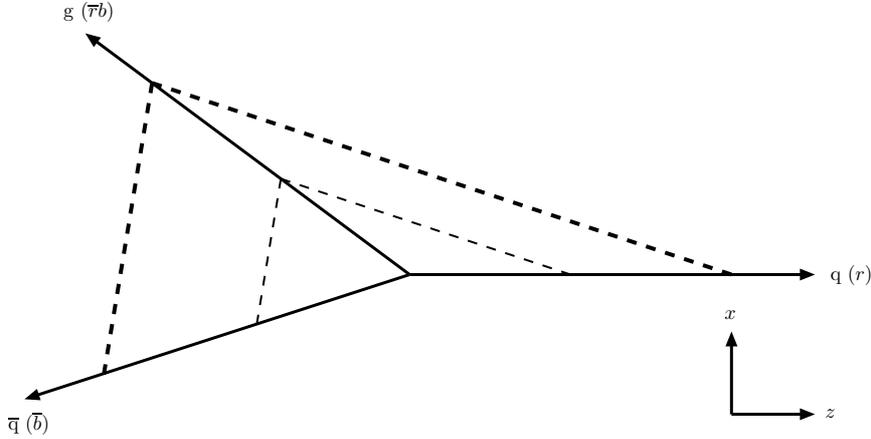


Figure 2: The string motion in a $q\bar{q}g$ system, neglecting hadronization. The q , \bar{q} and g move out from the common origin, all with the speed of light and along straight lines, in the limit that quark masses are neglected. The connecting dashed lines illustrate the string pieces at two times. A possible colour assignment is indicated within brackets.

the string drawing become more complicated. For a $q\bar{q}g$ event, a string is stretched from the q end via the g to the \bar{q} end, Fig. 2, 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 to quark string force is 2, a number which 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. Note that the factor 2 above is independent of the actual kinematical configuration: a smaller opening angle between two partons corresponds to a smaller string length being drawn out per unit time, but also to a larger transverse velocity of the string piece, thereby increasing its energy density. In fact, these two factors exactly cancel so that the energy loss per unit time indeed becomes the same.

The $q\bar{q}g$ string will fragment along its length. To first approximation this means that there is one fragmenting string piece between q and g and a second one between g and \bar{q} . One hadron is straddling both string pieces, i.e. sitting around the gluon corner. The rest of the particles are produced as in two simple $q\bar{q}$ strings, but strings boosted with respect to the overall CM frame. When considered in detail, the string motion and fragmentation is more complicated, with the appearance of additional string regions during the time evolution

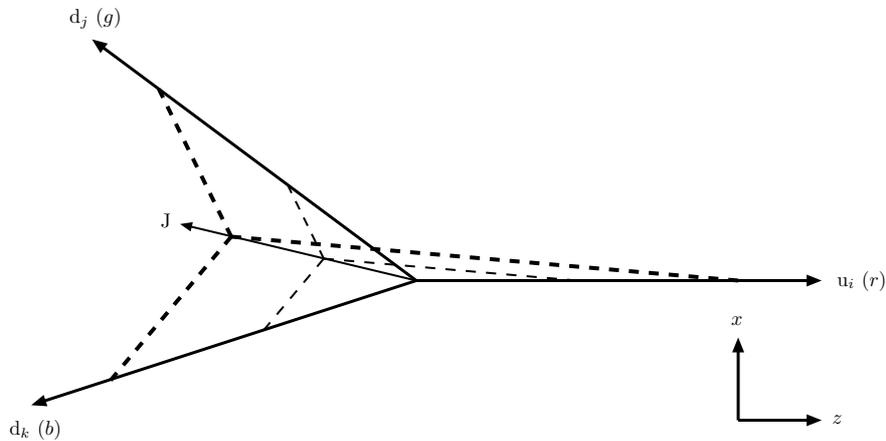


Figure 3: The string motion in a junction system, neglecting hadronization. The u_i , d_j and d_k move out from the common origin, all with the speed of light and along straight lines, in the limit that quark masses are neglected. The thin arrow indicates the resulting motion of the junction J . The connecting dashed lines illustrate the Y-shaped string topology at two discrete times. A possible colour assignment is indicated within brackets.

of the system [26]. These corrections are especially important for soft and collinear gluons, since they provide a smooth transition between events where such radiation took place and events where it did not. Therefore the string fragmentation scheme is “infrared safe” with respect to soft or collinear gluon emission.

4.2 The junction

As we saw above, meson states can be represented by little yo-yo’s, with a string piece stretched between the q and \bar{q} . What would be the corresponding representation for a baryon? In normal string fragmentation the issue is not so interesting, since the hadron size is consistent with the string width, meaning the internal structure is not really resolved. Thus the mesonic yo-yo is more of a convenient concept for book-keeping, and it is only for higher-mass systems that the elongated string structure really becomes of relevance. The equivalent situation for a baryon state would then be when several valence quarks are kicked out of an incoming proton or, as here, when baryon number is violated in Supersymmetry. In its simplest form, it can be illustrated by the decay $\tilde{\chi}_1^0 \rightarrow u_i d_j d_k$, or equivalently a proton in which all valence quarks are kicked out in different directions at the same time.

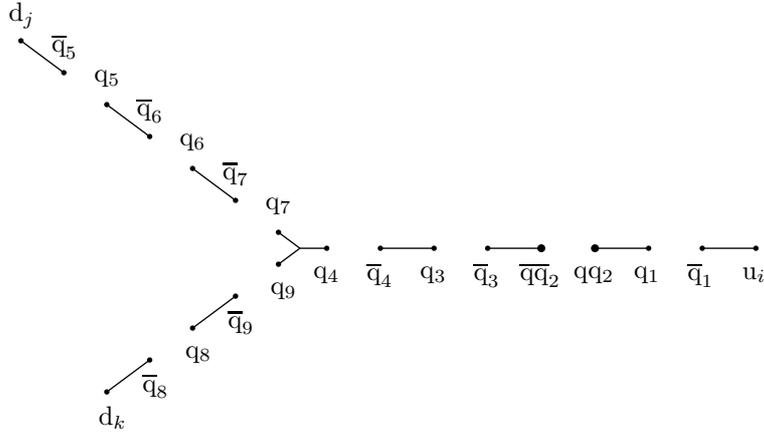


Figure 4: Hadronization by $q'\bar{q}'$ production in a junction string topology. The figure is in an abstract “flavour space”, to be related to a space–time topology like Fig. 3, with flavours belonging to the same hadron connected by string pieces. The labelling of new quarks is arbitrary. The $q_4q_7q_9$ hadron carries the original baryon number, while q_1q_2 and $\bar{q}_2\bar{q}_3$ correspond to the possibility of baryon–antibaryon pair production in the hadronization process, and the rest represents normal meson production.

Our solution here is to introduce the concept of a junction, in which a string piece from each of the three quarks meet, i.e. a Y-shaped topology, Fig. 3. Each of the three strings are of the normal kind already encountered. For instance, the string from the blue d_k quark acts as if there were an antiblue antiquark in the junction position. Recall that the colours of the two other quarks add like $g + r = \bar{b}$, given that the colour representation of the quarks is totally antisymmetric so as to make the state a colour singlet. Correspondingly for the other two string pieces. Only in the immediate neighbourhood of the junction would the field topology be more complicated.

In the hadronization, ordinary $q'\bar{q}'$ pair production will break up each of the three string pieces. This leads to meson production, and in diquark/popcorn scenarios also to baryon–antibaryon pair production, but not to the production of any net baryon number. Instead it is the three q' 's nearest to the junction, one from each string piece, that team up to form the baryon that is associated with the original baryon number, Fig. 4. In this sense the junction, by its colour topology, becomes the carrier of the baryon number of the system, rather than the three original quarks. This baryon typically comes to have a small momentum in the rest frame of the junction, as we shall see. This holds also

in the rest frame of a typical $\tilde{\chi}_1^0 \rightarrow u_i d_j d_k$ decay, where the three quarks move out in widely separated directions, since the junction then has a small velocity.

The concept of a junction is not new. It was introduced in the early string model of hadrons [27–29], and has been used in confinement studies to set up the colour structure of baryons in the same way as we do here [30]. Indeed, the simplest locally gauge invariant operator of baryon number unity that can be constructed in SU(3) is [29]:

$$B_{i_1 i_2 i_3} = \epsilon^{\alpha_1 \alpha_2 \alpha_3} \prod_{n=1}^3 P \left[e^{ig \int_{\mathcal{P}(x, x_n)} A_\mu dx^\mu} q_{i_n}(x_n) \right]_{\alpha_n} \quad (7)$$

where i_n and α_i are indices for flavour and colour, respectively, A_μ is the gluon field, and P expresses ordering along the path \mathcal{P} from x to x_n . Physically, this operator represents three quarks sitting at x_1 , x_2 , and x_3 , respectively, each connected via the gluon field to the point x , which we may identify as corresponding to the locus of the string junction. Thus, from this point of view, QCD does appear to select the junction topology.

The junction concept has also been applied to the understanding of the fragmentation function of baryons in low- p_\perp hadronic collisions [31], providing a more sophisticated variant of the popcorn mechanism whereby the two quarks of a diquark need not end up in the same hadron. Multiple perturbative parton–parton interactions have been suggested as a mechanism to activate the junction in proton beams [32]. And junctions have been proposed for $q\bar{q}g$ configurations, in scenarios where a separate kind of colour octet gluon string exists (which requires a string tension less than twice that of the normal colour triplet string, or else it would be energetically unfavourable) [33].

A simplified and unofficial implementation of junctions has existed in PYTHIA since many years, and been used e.g. from SUSYGEN. To the best of our knowledge, however, what we present here is the first complete scenario for the hadronization of generic junction topologies in high-energy interactions.

It could have been interesting to contrast the junction concept with some alternatives, but we have been unable to conceive of any realistic such, at least within a stringlike scenario of confinement. The closest we come is a V-shape topology, with two string pieces, similar to the $q\bar{q}g$ topology in Fig. 2. This would be obtained if one e.g. imagined splitting the colour of the upper left quark in Fig. 3 as $g = \bar{r}\bar{b}$. In such a scenario the baryon would be produced around this quark, and could be quite high-momentum. Of course, such a procedure is arbitrary, since one could equally well pick either of the three quarks to be in the privileged position of producing the key baryon. Further, with two string pieces now being pulled out from one of the quarks, the net energy stored in the string at a given (early) time is larger than in the junction case, meaning the Y junction is energetically favoured over the V topology. For these reasons, the V scenario has not been pursued.

It should be noted that, in our junction scenario, it can happen that one quark has a low initial momentum in the junction rest frame, and therefore never strays very far away. The string between this quark and the junction need then not break, so that the quark ends up in the central baryon together with the junction, i.e. the quark and the junction effectively act as a unit. Seemingly this is akin to the scenario sketched above but, in the junction topology framework, it could only happen when the quark has low momentum in the junction rest frame, so again the key baryon would be low-momentum.

Other alternatives appear even more far-fetched. For instance, we see no possibility for a Δ -shaped topology, i.e. connecting each quark directly with the other two. Not only would it be even more energetically unfavourable relative to the junction topology, but it is also difficult to conceive of a colour assignment that would at all allow it. One would then need to introduce a novel kind of strings, that do not obey the normal QCD colour charge and flux quantization conditions.

4.3 Junction motion

In the rest frame of the junction the opening angle between any pair of quarks is 120° , i.e. we have a perfect Mercedes topology. This can be derived from the action of the classical string [27], but follows more directly from symmetry arguments. What is maybe not so obvious is that the 120° condition also is fulfilled if the quarks are massive. However, in the junction rest frame each string piece is pulled straight out from the origin, without any transverse motion, so the string pieces do not carry any momentum. Then overall momentum conservation gives that

$$\frac{d\mathbf{p}_{\text{tot}}}{dt} = \sum_{\text{q}} \frac{d\mathbf{p}_{\text{q}}}{dt} = -\kappa \sum_{\text{q}} \mathbf{e}_{\text{q}} = \mathbf{0} . \quad (8)$$

Here we have used eq. (4) to give the momentum loss of quarks of arbitrary mass, where \mathbf{e}_{q} is a unit vector along the direction of motion of each of the three quarks.

The flow of energy and momentum is akin to that in the normal $q\bar{q}$ system, where the intermediate string absorbs and cancels the momentum of the receding q and \bar{q} . With respect to the qqq case, it is implicit in eq. (8) that no net momentum is accumulated by the junction. Instead, it acts as a switchyard for the momentum flowing in the system, thereby cancelling the momentum given up by the three endpoint quarks. In this way, the junction itself remains with a vanishing four-momentum, which obviously holds in an arbitrary reference frame. In a general frame the string pieces would have a transverse motion as well, however, and thereby carry a nonvanishing momentum. Then eq. (8) would need to be generalized to include these additional terms and would become less transparent.

The rest frame of the junction can easily be found for the case of three massless quarks (and no further gluons), but the general massive case admits no analytical solution. A convenient numerical solution can be obtained as follows. Imagine three quarks with four-momenta p_1, p_2 and p_3 . The four-products $a_{ij} = p_i p_j$ are Lorentz invariants, and thus the boosted momenta p'_i obey

$$a_{ij} = p'_i p'_j = E'_i E'_j - \mathbf{p}'_i \mathbf{p}'_j = E'_i E'_j - |\mathbf{p}'_i| |\mathbf{p}'_j| \cos \frac{2\pi}{3} = E'_i E'_j + \frac{1}{2} |\mathbf{p}'_i| |\mathbf{p}'_j| \quad (9)$$

in the junction rest frame. This leads us to introduce

$$f_{ij} \equiv f(|\mathbf{p}'_i|, |\mathbf{p}'_j|; m_i, m_j, a_{ij}) = \sqrt{|\mathbf{p}'_i|^2 + m_i^2} \sqrt{|\mathbf{p}'_j|^2 + m_j^2} + \frac{1}{2} |\mathbf{p}'_i| |\mathbf{p}'_j| - a_{ij} . \quad (10)$$

Note that f_{ij} is a monotonically increasing function of each of its first two arguments. If we choose e.g. to let $|\mathbf{p}'_1|$ vary freely within the kinematically allowed region, the requirements $f_{12} = 0$ and $f_{13} = 0$ can then be uniquely solved to give the other two momenta,

$$|\mathbf{p}'_j| = \frac{2E'_1 \sqrt{4a_{1j}^2 - m_j^2 (4E_1'^2 - |\mathbf{p}'_1|^2)} - 2|\mathbf{p}'_1| a_{1j}}{4E_1'^2 - |\mathbf{p}'_1|^2} , \quad (11)$$

which both decrease with increasing $|\mathbf{p}'_1|$. Therefore also f_{23} is monotonically decreasing if viewed as a function of $|\mathbf{p}'_1|$, $f_{23} = f(|\mathbf{p}'_2|(|\mathbf{p}'_1|), |\mathbf{p}'_3|(|\mathbf{p}'_1|))$. The final requirement $f_{23} = 0$ now gives a unique solution $|\mathbf{p}'_1|$. This solution can be obtained by an iterative interpolating procedure. In the massless case, the solution simplifies to $E'_i = |\mathbf{p}'_i| = \sqrt{2a_{ij} a_{ik} / 3a_{jk}}$.

Once the energies in the junction rest frame are known, the construction of a boost is straightforward. A first boost with $\beta^{\text{CM}} = -\sum \mathbf{p}_i / \sum E_i$ brings the three quarks to their rest frame, with known momenta p_i^{CM} . A second boost β' , that brings the system to the junction rest frame, now obeys the three equations $\gamma' E_i^{\text{CM}} + \gamma' \beta' \mathbf{p}_i^{\text{CM}} = E'_i$. After division by E_i^{CM} , and subtraction of two of the equations from the third, one obtains e.g.

$$\gamma' \beta' \left(\frac{\mathbf{p}_1^{\text{CM}}}{E_1^{\text{CM}}} - \frac{\mathbf{p}_j^{\text{CM}}}{E_j^{\text{CM}}} \right) = \frac{E'_1}{E_1^{\text{CM}}} - \frac{E'_j}{E_j^{\text{CM}}} . \quad (12)$$

Since the three vectors \mathbf{p}_i^{CM} lie in a plane, the boost vector can be parameterized as a linear sum of the two difference vectors defined by the above equation for $j = 2$ and 3 , which gives a simple linear equation system for $\gamma' \beta'$. Finally, the overall boost is obtained by combining β^{CM} with β' .

In the general case, there does not have to exist a solution with perfect symmetry. Such troublesome events, characterized e.g. by $f_{23} < 0$ for all $|\mathbf{p}'_1|$,

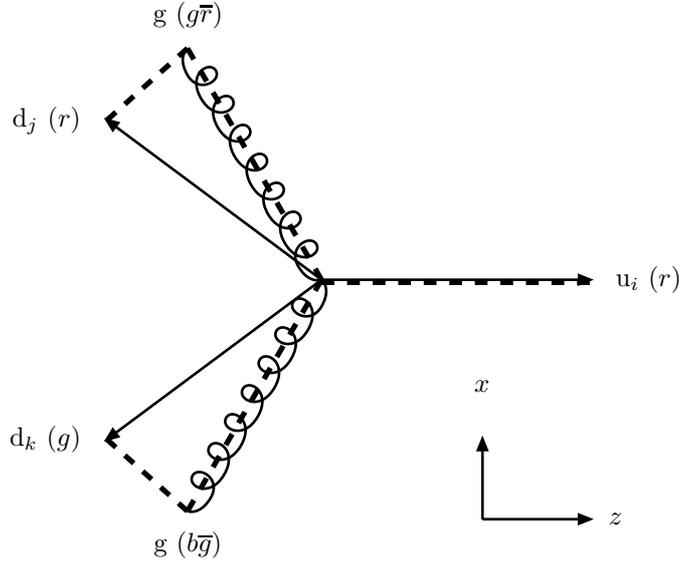


Figure 5: String topology of a $\tilde{\chi}_1^0 \rightarrow u_i d_j d_k$ decay, where the two $d_{(j,k)}$ quarks each has radiated a gluon. The event is drawn in the rest frame of the junction at early times, where the u and the two gluons are separated by 120° , and the string topology is shown by dashed lines. A possible colour assignment is indicated within brackets.

are very rare, however. We have only encountered them when a massive quark is almost at rest in the frame that comes closest to giving a symmetrical topology. We here accept an imperfect solution somewhere in the neighbourhood of this singular point. Since it is the size of the boost that will matter in the hadronization description, not the 120° opening angles *per se*, this should be acceptable.

So far we have assumed that the junction remains in uniform motion. When gluon emission is included, this need no longer be the case. Consider e.g. an event like the one in Fig. 5. Here the two $d_{(j,k)}$ quarks each radiated a gluon, and so the strings from these quarks to the junction are drawn via the respective gluon, cf. Fig. 2. It is the direction of these gluons, together with the u quark, that determines the junction motion at early times, and the directions of the $d_{(j,k)}$ quarks themselves are irrelevant. As a gluon moves out from the junction origin, it loses energy at a rate $dE/dt = -2\kappa$, where the 2 comes from it being attached to two string pieces. So after a time $E_g(0)/2\kappa$ it has lost all its original energy, and after a further equal time this information has propagated back to

the junction. From then on it would be the direction of the respective $d_{(j,k)}$ quark, and not of the gluon, that defines the pull on the junction. In the example of Fig. 5, the junction would originally be at rest but later on start to move leftwards. In the generic configuration, each of the outgoing quarks could radiate several gluons, that would be colour-connected in a chain stretching from the endpoint quark inwards to the junction. The innermost gluon on each of the three chains would together determine the original motion, with the next-innermost taking over after the innermost lost its energy, and so on, spreading outwards. As a consequence the junction would “jitter around”.

Now, all of this is moot if the string starts to fragment before the innermost gluon lost its energy, because a string break will sever the inward flow of momentum. As we noted above, the typical scale for this to happen is of order $\langle \kappa\tau \rangle \approx 1.5$ GeV. Therefore an innermost gluon with energy much above this scale by itself defines the junction motion, whereas a gluon with energy much below would act on the junction too short a time to matter.

Rather than trying to trace the junction jitter in detail — which anyway will be at or below the limit of what it is quantum mechanically meaningful to speak about — we define an effective pull of each string on the junction as if from a single particle with a four-momentum

$$p_{\text{pull}} = \sum_{i=1}^n p_i \exp\left(-\sum_{j=1}^{i-1} E_j/E_{\text{norm}}\right) . \quad (13)$$

Here $i = 1$ is the innermost gluon, $i = 2$ is the next-innermost one, and so on till the endpoint quark $i = n$. The energy sum in the exponent runs over all gluons inside the one considered (meaning it vanishes for $i = 1$), and is normalized to a free parameter E_{norm} , which by default we associate with the $\langle \kappa\tau \rangle$ above. Note that the energies E_j depend on the choice of frame. A priori, it is the energies in the rest frame of the junction which should be used in this sum, yet since these are not known to begin with, we use an iterative procedure, starting with the (known) energies in the CM of the string system, calculating the corresponding pull vectors in this frame and from them the candidate junction rest frame, calculating the pull vectors in this new frame and so forth until the Lorentz factor of the last boost is $\gamma_{\text{last}} < 1.01\gamma_{\text{tot}}$.

4.4 Junction hadronization

As we have noted above, hadronization begins in the middle of an event and spreads outwards. In the junction rest frame the junction baryon would thus be the one to form first, on the average, and have a rather small momentum. Thereafter, each of the three strings would fragment pretty much as ordinary strings, e.g. as in a back-to-back $q\bar{q}$ pair of jets. Also extensions to systems with multiple gluon emissions should closely follow the corresponding pattern for normal events. With or without gluon emission, we shall speak of three

strings connected at the junction, where a string may consist of several string pieces between adjacent (i.e. colour-connected) partons.

In particular, if we consider events where each of the three outgoing quark jets have large energies in the junction rest frame, the production of high-momentum particles inside a jet should agree with the one of a corresponding jet in an ordinary two-jet event. This can be ensured by performing the fragmentation from the outer end of the strings inwards, just like for the $q\bar{q}$ string. Thus an iterative procedure can be used, whereby the leading q is combined with a newly produced \bar{q}_1 , to form a meson and leave behind a remainder-jet q_1 , which is fragmented in its turn. Flavour rules, fragmentation functions and handling of gluon-emission-induced kinks on the string are identical with the ones of the ordinary string.

While these hadronization principles as such are clear, and give the bulk of the physics, unfortunately there is a catch: if all three strings are fragmented until only little energy and momentum remain in each, and then these remainders are combined to a central baryon, what guarantees that this baryon obtains the correct invariant mass it should have?

The same problem exists in simple $q\bar{q}$ events, that also there energy and momentum is not guaranteed to come out right, if the event is fragmented independently from the two ends and then joined by a single hadron in the middle. In this case, the overall energy-momentum conservation is obtained by the “area law” [21], which couples the production of *all* the hadrons in the event. It is therefore difficult to simulate exactly, but for large remaining invariant masses it simplifies to the iterative framework already introduced. For small masses a pragmatic approximation is used. Firstly, not one but two hadrons are used to join the jets. A two-hadron system has a continuous mass spectrum, so one may construct consistent kinematics if the normal fragmentation is stopped when the remaining invariant mass has dropped below some value. Since the final two hadron masses are not known beforehand, sometimes the remaining mass drops below the two-body threshold, in which case the fragmentation is re-started from the beginning. Secondly, by a random choice of producing the next hadron either off the q end or off the \bar{q} one, the final joining does not always occur exactly in the middle of events, thereby smearing remaining imperfections of the joining procedure proper.

For the fragmentation of a junction topology, we attempt to retain as much as possible the known good features of the existing approach to $q\bar{q}$ events, although this involves conflicting interests, as follows. In order to describe the production of high-momentum particles, fragmentation should be allowed to proceed from all three string ends inward. But, in order not to bias the junction baryon overly, the joining for energy-momentum conservation should not always have to influence this hadron. Briefly put, the solution is to fragment two of the three strings inwards, thereafter combine their leftovers to an effective diquark, and finally fragment the string between this diquark and the third end

in much the same spirit as described for $q\bar{q}$ events above, i.e. by fragmentation at random off both ends of the system. Put in the context of Fig. 4, imagine first tracing the chains $d_j - q_5 - q_6 - q_7$ and $d_k - q_8 - q_9$, then forming a diquark q_7q_9 , and thereafter fragmenting the string between this diquark and u from both ends.

There are a number of technical details, as follows. The hadronization process itself is normally carried out in the rest frame of the colour singlet system under consideration, but is Lorentz covariant, so another choice would be no problem. Information on the junction motion is encoded in its velocity four-vector $v_{\text{jun}} = (\gamma', -\gamma'\beta')$, with the β' defined in subsection 4.3.

For normal string pieces, spanned between two partons, the invariant mass of the piece is defined by the four-momenta of the two endpoint partons. Since the junction does not carry a momentum, there is no corresponding definition for the string pieces spanned between the junction and each of its three nearest colour-connected partons. Instead the junction end is here represented by a fictitious parton, specific to each string, opposite to the p_{pull} vector of eq. (13), as viewed in the junction rest frame, which gives $p_{\text{opp}} = -p_{\text{pull}} + 2(v_{\text{jun}}p_{\text{pull}})v_{\text{jun}}$.

Two of the three strings are fragmented from the respective end inwards, towards a fictitious other end as defined above. In order to have a large-mass system left for the system in which the joining occurs, we prefer to pick these two to be the ones with lowest energy, as defined in the junction rest frame, i.e. with lowest $E'_{\text{str}} = v_{\text{jun}} \sum_i p_i^{\text{CM}}$. Here p_i are the four-momenta of the partons belonging to the given string, excluding the fictitious junction one. As the hadrons are successively produced, their summed energy E'_{had} (in the same frame) is also updated. Once the hadronic energy exceeds the string one, $E'_{\text{had}} > E'_{\text{str}}$, the process has gone too far, i.e. passed the junction point of the string system, so it is stopped and the latest hadron is rejected.

The random choices made in the fragmentation function allows the energy in this latest hadron to fluctuate significantly. It can therefore be that, after its removal, the energy $\delta E' = E'_{\text{str}} - E'_{\text{had}}$ remaining of a string can be quite significant. This is particularly dangerous if it happens in both of the strings considered, since then the leftovers would be combined to give a momentum intermediate in direction to the two strings, and thereby maybe give a jet where none originally existed. Therefore the two hadronic chains are rejected and new ones are generated in their place if both $\delta E'$ are larger than a parameter $\delta E'_{\text{min}}$, by default 1 GeV. One of the two can have a larger energy than this, since then the combined leftovers would still essentially agree with the direction of this original string, but the chains are also rejected if this one has an energy above $\delta E'_{\text{min}} + R\delta E'_{\text{max}}$, where $\delta E'_{\text{max}}$ by default is 10 GeV and R is a random number uniformly selected between 0 and 1. The two parameters $\delta E'_{\text{min}}$ and $\delta E'_{\text{max}}$ are “tuned” by the requirements of a consistent description, see below. In order to avoid infinite loops, at most 10 attempts of the above kind are rejected.

When two acceptable hadronic chains have been found, the remaining four-

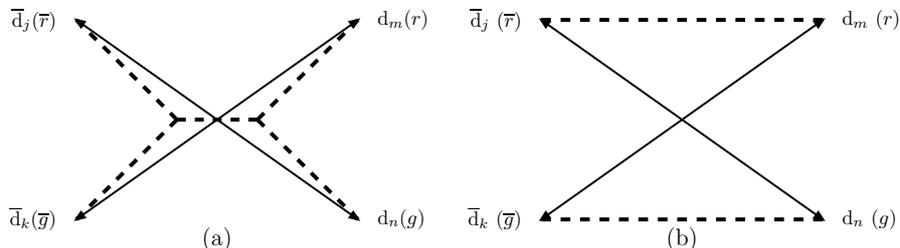


Figure 6: (a) Schematic string topology of a $e^+e^- \rightarrow \tilde{t}\tilde{t}^* \rightarrow \bar{d}_j\bar{d}_k d_m d_n$ decay. The strings are shown as dashed lines. A possible colour assignment is indicated within brackets. The string between the two junctions would then be blue–antiblue. (b) Alternative string drawing for the same process, without any junctions.

momenta from the respective two strings are combined into a single parton, which then replaces the junction as endpoint for the third string. (Actually, technically, the whole of the two string four-momenta is assigned to a parton, but then the fragmentation process is assumed already to have produced the given set of hadrons. This is almost equivalent, apart from some minor details of transverse momentum handling.) The new effective parton may have a larger momentum than energy, and thereby nominally be spacelike. If only by a little, it normally would not matter, but in extreme cases the whole final string may even come to have a negative squared mass. Therefore additional checks are made to ensure that the final string mass is above the threshold for string fragmentation. If so, the fragmentation procedure is identical with that of an ordinary string from here on, else repeated attempts are made, starting over with the first two strings.

A further aspect is the flavour properties. Except at the junction, the normal rules of string fragmentation are used. At the junction, the two strings fragmented first each define a leftover quark, which are combined into a diquark. This diquark is assigned a spin 0 or 1 according to the same relative probabilities as normally. Since also ordinary fragmentation can produce diquark–antidiquark pairs, it could be that either string ends with a leftover antidiquark rather than quark. Such configurations are not acceptable, and have to be rejected in favour of new tries. Once a diquark is defined, this can fragment further as usual. In the simplest scenario, it will next fragment to produce a baryon and a leftover antiquark. If popcorn baryon production is allowed, however, a meson may alternatively be split off to produce a new diquark. That is, the baryon number may then migrate to higher energies than otherwise, but will still be rather centrally produced.

4.5 Other string topologies

So far we have considered the production of string systems containing a single junction, well illustrated by the case of neutralino decays, but also possible in other processes. There is a second possibility, however, namely that of a system consisting of two junctions, or, more precisely, of a junction and an antijunction, where the former is associated with a baryon number +1 and the latter -1. An example would be $e^+e^- \rightarrow \tilde{t}\tilde{t}^* \rightarrow \bar{d}_j\bar{d}_k d_m d_n$, containing two BNV decays, $\tilde{t} \rightarrow \bar{d}_j\bar{d}_k$ and $\tilde{t}^* \rightarrow d_m d_n$. This kind of topology is illustrated in Fig. 6a. There are now two quark ends, two antiquark ones, and five strings, including the one between the two junctions. Each of these strings could contain an arbitrary number of intermediate gluons, where gluons in between the two junctions are related to emission in the production process $e^+e^- \rightarrow \tilde{t}\tilde{t}^*$ and the others to the decay processes (with the standard ambiguities, especially relevant for the soft gluons).

The junction motion and string fragmentation follows by simple generalization of what has already been discussed for configurations with a single junction. Like before, two of the three strings from a junction — the two not connected to the other junction — are first fragmented, and the leftovers combined to an effective (anti)diquark. When this has been performed for both junctions, what remains is a diquark–antidiquark string spanned between the two junctions, which then is fragmented. Thus the hadronization will produce one baryon and one antibaryon, associated with the junctions, plus potentially additional baryon–antibaryon pairs as usual.

The same checks as above have to be applied at various stages, but no new ones. The one ambiguity is how to calculate the p_{pull} of eq. (13) for the third string of a junction, hooking up to the other junction. Here gluons on the string between the two junctions are considered as normally, while partons on the far side of the other junction are tracked down each of the two strings without including any exponential suppression from energies in the other of the strings.

Since the net baryon number vanishes in the above topology, an alternative string scenario is illustrated by Fig. 6b. Here two strings are pulled directly between a $d_{(m,n)}$ and a $\bar{d}_{(j,k)}$ quark. The junctions are gone, and so such an event would not have to contain any baryons or antibaryons at all, and definitely not any associable with the BNV processes as such. This is the approach adopted in HERWIG.

A priori, one does not know which of the two above scenarios would be the correct one. It would not even have to be a unique answer, valid for all events. A possibility would be that the topology with minimal string length is selected dynamically. Then events of the general kind shown in Fig. 6 would obtain the 2-junction topology of Fig. 6a when the $d_m d_n$ and $\bar{d}_j \bar{d}_k$ opening angles are small, while the 0-junction one of Fig. 6b would result if instead the $d_m \bar{d}_j$ and $d_n \bar{d}_k$ opening angles were the small ones.

Such a simple picture would become more complicated by the addition of

gluons in the production process, i.e. the ones we above put on the string between the two junctions. It now becomes necessary to subdivide such gluons into two separate colour lines, one between each of the outgoing $d_{(m,n)}\bar{d}_{(j,k)}$ pairs. In general, we would expect the 0-junction topology to become more disfavoured when there are several gluons being emitted.

Furthermore, in the limit of large \tilde{t} lifetime, the string between the \tilde{t} and \tilde{t}^* would have time to start to fragment before the \tilde{t} or \tilde{t}^* has decayed. In such a case, the event would consist of two separate colour singlets, that fragment independently of each other. This does not have to mean timescales long enough to have multiple string breaks and the formation of stop-hadrons: one single string break is enough to make the production of a baryon and an antibaryon unavoidable. At first glance a sufficiently large lifetime for this to happen, $c\tau \gtrsim 1$ fm, would seem unlikely. However, when only BNV decay channels are kinematically allowed, the \tilde{t} lifetime is almost always non-negligible, since, for massless d, s, and b quarks the decay length is roughly:

$$c\tau_{\tilde{t}}^{\text{BNV}} \approx (100 \text{ fm}) \times \left(\frac{0.01}{|\lambda_3''|} \right)^2 \times \frac{100 \text{ GeV}}{m_{\tilde{t}}} \quad (14)$$

where $|\lambda_3''|$ represents the average of the $|\lambda_{3jk}''|$ couplings. Strictly speaking, the above formula is valid for \tilde{t}_R . For \tilde{t}_1 (\tilde{t}_2), it is increased by $1/|\sin\theta_{\tilde{t}}|^2$ ($1/|\cos\theta_{\tilde{t}}|^2$). In passing, we note that event properties are likely to be similar whether stop-hadrons have time to be produced or not: corresponding studies for ordinary top quarks have shown that differences are restricted to the region of low-momentum particles, and there mainly show up in angular distributions, not averaged event properties [34].

If gauge-interaction decay channels such as $\tilde{t} \rightarrow b\tilde{\chi}_1^+$ are open, these will almost always dominate. Then, ignoring what happens to the colour disconnected $\tilde{\chi}_1^+$, it would be doubly rare to have both \tilde{t} and \tilde{t}^* decay with BNV. Experimental studies would presumably have to concentrate on finding events where one decay is BNV and the other not, and so we are back with a 1-junction topology.

Finally, we note that $\tilde{t}\tilde{t}^*$ production in hadron colliders would predominantly come from $gg \rightarrow \tilde{t}\tilde{t}^*$ in the colour octet channel, where thereby the \tilde{t} and \tilde{t}^* would belong to separate colour singlet subsystems. Again the production of a baryon and an antibaryon would then be inevitable.

Nevertheless, in order to test the consequences, we have developed an alternative model. For a given event with two colour-connected BNV decays, the total string length is evaluated under the assumption that the event either is arranged in a 2-junction or in a 0-junction topology, and the topology is chosen that corresponds to the smallest length. Given this choice the subsequent hadronization is well-defined.

In more detail, the string length is defined by the so-called λ measure [35].

For a simple two-parton system

$$\lambda = \ln \left(\frac{s}{m_0^2} \right) . \quad (15)$$

Here s is the squared invariant mass of the system and m_0 is some hadronization reference scale of order $\langle \kappa\tau \rangle$, just like the E_{norm} defined in section 4.3. Particles are produced with a flat rapidity distribution in the central regions of the string, while the distribution falls off near the ends. Then $y_{\text{max}} \simeq \ln(\sqrt{s}/m_0)$ defines the effective rapidity range, in the middle of the fall-off, such that the total multiplicity on the average is proportional to y_{max} . For a generic string configuration, with many gluons between the quark and antiquark, the complete λ expression is rather messy, but approximately it can be represented by the linear sum of the λ measure for each string piece,

$$\lambda = \sum_{i=1}^{n-1} \lambda_{i,i+1} = \sum_{i=1}^{n-1} \ln \left(\frac{s_{i,i+1}}{m_0^2} \right) . \quad (16)$$

Here the squared invariant mass $s_{i,i+1}$ is calculated with the full momentum for the endpoint quarks but only half for the intermediate gluons, which are shared between two adjacent string pieces.

We now need to generalize the λ measure to the case with junctions, which has not been considered in the literature so far. To begin, revert to a simple back-to-back $q\bar{q}$ system, with quark masses negligible. Then

$$\begin{aligned} \lambda &= \ln \left(\frac{s}{m_0^2} \right) \\ &= \ln \left(\frac{4E_q E_{\bar{q}}}{m_0^2} \right) \\ &= \ln \left(\frac{2E_q}{m_0} \right) + \ln \left(\frac{2E_{\bar{q}}}{m_0} \right) \\ &= \ln \left(\frac{2vp_q}{m_0} \right) + \ln \left(\frac{2vp_{\bar{q}}}{m_0} \right) . \end{aligned} \quad (17)$$

The splitting into two terms can be seen as a separation of the full rapidity range into one on the quark side of the event and another on the antiquark one, smoothly matching at the origin. At this stage the origin is arbitrary, i.e. $E_q \neq E_{\bar{q}}$ represents an event boosted along the event axis; the individual terms may thus be changed but the sum remains invariant. In the final step the origin is represented by the four-vector $v = (1; 0, 0, 0)$, to allow a Lorentz invariant form also for the split expression. Alternatively to considering $E_q \neq E_{\bar{q}}$ one may then use a v boosted along the event axis.

In a $q_1q_2q_3$ event, if viewed in the rest frame of the junction, there will be three rapidity plateaus extending from the origin, so one can calculate a total

rapidity range

$$\lambda = \ln \left(\frac{2vp_1}{m_0} \right) + \ln \left(\frac{2vp_2}{m_0} \right) + \ln \left(\frac{2vp_3}{m_0} \right). \quad (18)$$

The v four-vector now is easily identified with the motion of the junction.

We also need to consider the rapidity length of the string between two junctions. Since the rapidity plateau extends all the way to the junction, this is actually given by the rapidity difference between the junctions themselves. Evaluating this in a frame where the junctions are back-to-back, $v_{1,2} = (\gamma; 0, 0, \pm\gamma\beta)$, one obtains

$$\lambda = 2y_{\max} = 2\frac{1}{2} \ln \left(\frac{\gamma + \gamma\beta}{\gamma - \gamma\beta} \right) = \ln \left(\frac{1 + \beta}{1 - \beta} \right) = \ln \left(\frac{1 + \beta^2 + 2\beta}{1 - \beta^2} \right), \quad (19)$$

with $v_1v_2 = \gamma^2(1 + \beta^2) = (1 + \beta^2)/(1 - \beta^2)$ or $\beta^2 = (v_1v_2 - 1)/(v_1v_2 + 1)$.

We can now give the λ of the two configurations in Fig. 6a and 6b. Since the number of m_0 factors agrees between the two, it is more convenient to introduce $\Lambda = m_0^4 \exp(\lambda)$, with expressions

$$\Lambda_{2\text{-junction}} = (2v_1p_m)(2v_1p_n)(2v_2p_j)(2v_2p_k) \left(v_1v_2 + \sqrt{(v_1v_2)^2 - 1} \right) \quad (20)$$

$$\Lambda_{0\text{-junction}} = (2p_m p_j)(2p_n p_k), \quad (21)$$

where we have reused the generation indices of the quarks to distinguish them, and v_1 (v_2) corresponds to the right (left) vertex. The smaller of the two Λ values now determines which configuration would be the preferred one. Note that both expressions are linear in each of the p_i , so the choice only depends on the directions of motion of the outgoing quarks.

Now consider the complications caused by showering. Radiation after the BNV decays is no problem: one only needs consider the connection out to the colour produced in the BNV decays, i.e. to the gluon nearest to the junction, since the subsequent colour chain out to the endpoint quark would be unaffected by the colour arrangement between the two BNV vertices. The $p_{j,k,m,n}$ in eqs. (20)–(21) should be redefined accordingly. Radiation patterns before the BNV decays can be divided into two classes.

(i) If the shower contains a $g \rightarrow q\bar{q}$ branching, the colour flow in the system is automatically broken, and then the production of a baryon–antibaryon pair is unavoidable. This is the perturbative equivalent of the string between \tilde{t} and \tilde{t}^* starting to fragment before the \tilde{t} or \tilde{t}^* has decayed. As one moves up from the kinematical threshold this probability increases. Since the collinear and soft gluon singularities are tamed by the large \tilde{t} mass and width, the total $g \rightarrow q\bar{q}$ probability is reasonably reliably predictable, although a cut-off dependence e.g. from the assumed effective quark masses remains.

(ii) If there are no $g \rightarrow q\bar{q}$ branchings, all the gluons in the shower can be

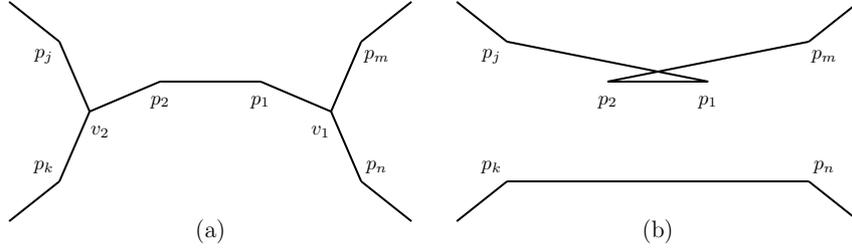


Figure 7: Illustration of (a) the 2-junction topology and (b) a 0-junction topology, for identical parton configurations. Indices have been chosen so as to correspond directly to the expressions in eqs. (22) and (23). Naturally, the topology in (b) only corresponds to one possible choice of which is which of (m,n) and (j,k) .

arranged in a single colour chain between the \tilde{t} and \tilde{t}^* . When the BNV decays occur, the anticolour of the gluon that matches the colour of the \tilde{t} could no longer match the anticolour of either \tilde{t} decay product in a junction-free topology, but would rather have to be matched to the colour of either \tilde{t}^* decay product. Thus the 0-junction configuration would be with two strings, one between one of the \tilde{t}^* decay product traversing all the gluons from the original \tilde{t} edge to the \tilde{t}^* one and then zig-zagging back to one of the \tilde{t} decay products, while the other string would be between the two remaining decay products. See Fig. 7 for an illustration of (a) the 2-junction configuration and (b) the 0-junction one just discussed. (Actually, since the stop mass reduces collinear emission, and since the stop decay products are widely scattered, the zig-zag pattern need not be as extreme as one might guess at first glance.) Colour-suppressed topologies are neglected, such as where the gluons split off into a closed loop and then two strings are stretched directly between the \tilde{t} and \tilde{t}^* decay products. When comparing string length with and without junctions, the string internally between the intermediate gluons is common and can be neglected. Labelling by p_i the gluon closest to the v_i vertex in colour space, in the chain between the two junctions, and now putting $\Lambda = m_0^6 \exp(\lambda)$, one obtains

$$\Lambda_{2\text{-junction}} = (2v_1 p_m)(2v_1 p_n)(2v_1 p_1)(2v_2 p_j)(2v_2 p_k)(2v_2 p_2), \quad (22)$$

$$\Lambda_{0\text{-junction}} = (2p_m p_2)(2p_j p_1)(2p_n p_k), \quad (23)$$

where which is which of m and n (and of j and k) is chosen at random.

If anything, the approach above could overestimate the probability of no baryon–antibaryon production, since it only considers the gain in string length once one has arrived at a new colour arrangement, whereas there could be

additional dynamic suppressions on the way between the ‘original’ 2-junction topology and the ‘final’ 0-junction one. Nevertheless, it should offer some realistic estimates how big a loss of baryon signal could at worst be. To quantify this to some extent, we have taken a closer look at isolated decays of $\tilde{t}\tilde{t}^*$ pairs (colour connected to each other to form an overall colour singlet) in the CM of the pair.

For high-momentum stops, $\gamma_{\tilde{t}} \equiv \frac{E_{\text{CM}}}{2m_{\tilde{t}}} \gg 1$, the $d_m d_n$ pair lies at a large rapidity separation from the $\bar{d}_j \bar{d}_k$ pair. In the 2-junction topology, this rapidity range is spanned by *one* string piece, the junction–junction one, whereas, in the 0-junction topology, *each* of the $d\bar{d}$ string pieces must cross it. In fact, due to the see-saw nature of the 0-junction topology in the presence of gluon radiation from the stops, the central rapidity range will be crossed more than twice if the stops radiate. Thus, we expect the 2-junction topology to dominate for cases $\gamma_{\tilde{t}} \gg 1$.

On the other hand, when the stops are slow, $\gamma_{\tilde{t}} \sim 1$, the \tilde{t} and \tilde{t}^* decay products are widely spread by the decay kinematics. Typically the invariant mass between a \tilde{t} decay product and a \tilde{t}^* one will be smaller than that between the two products of the same \tilde{t} or \tilde{t}^* , such that the 0-junction topology is guaranteed to give a smaller string length. Also for somewhat larger $\gamma_{\tilde{t}}$, where one can start to speak of an event axis, decay products are often thrown into the opposite hemisphere, such that the central rapidity region need not be crossed at all in the 0-junction alternative. Thus, we expect the 0-junction topology to dominate for reasonably low $\gamma_{\tilde{t}}$.

The rapidity arguments presented in the above two paragraphs closely follow the here defined string length measures, eqs. (20)–(23). However, we may alternatively use an argument based on the string motion itself. Recall that a junction is at rest when the opening angles between its three attached strings is 120° . Thus, if the opening angles between the motion of the \tilde{t} and its two decay products is 60° each, and correspondingly for the \tilde{t}^* , we are exactly at a balancing point. For stops faster than this, the smaller decay angles will force the two junctions to move away from each other, and so it should be less likely for them to meet and annihilate. For slower stops, the junctions actually move *towards* each other (if assumed produced at a distance apart), making it more believable that they can annihilate. Although these deliberations are separate from the string length minimization procedure described above, and do not give the same answer event by event, they do corroborate it by leading to the same expected dominant topologies as $\gamma_{\tilde{t}}$ is varied. The two approaches are about equally credible, but the string length argument has the advantage of leading to tractable answers also for quite complicated topologies.

Using now eqs. (20)–(23) to calculate total string lengths for two different stop masses, and a range of stop boosts $1 \leq \gamma_{\tilde{t}} \leq 6$, gives the plots shown in Fig. 8. As is apparent from the discussion above, we expect the ratio of 2-junction to 0-junction string topologies to depend mainly on $\gamma_{\tilde{t}}$ and only weakly

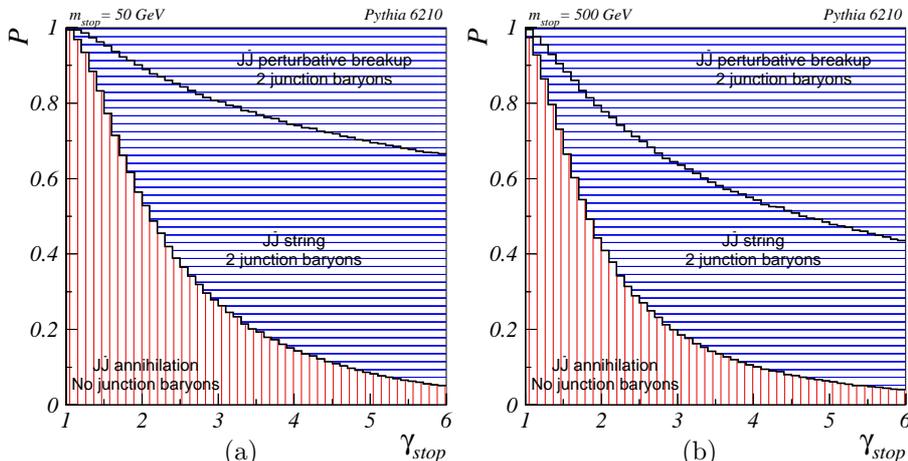


Figure 8: Probabilities for 2-junction and 0-junction configurations as functions of $\gamma_{\tilde{t}} = \frac{E_{CM}}{2m_{\tilde{t}}}$ for (a) $m_{\tilde{t}} = 50 \text{ GeV}$ and (b) $m_{\tilde{t}} = 500 \text{ GeV}$.

on the stop mass as such, and this agrees with what is shown in the relation between the lower two regions of the plots. The feature which *does* change significantly is the rate of perturbative breakups of the 2-junction systems, i.e. $g \rightarrow q\bar{q}$ splittings in the parton shower initiated before the BNV decays, point (i) above. The $g \rightarrow q\bar{q}$ rate increases with larger total energy, also for fixed $\gamma_{\tilde{t}}$: while the primary emission rate of gluons off the $\tilde{t}\tilde{t}^*$ system is rather constant, the phase space for further cascading of these gluons is increased with energy.

5 Model tests

In this section we will present some distributions that illustrate the basic properties of our model, and show that it works as expected. In several cases we will compare with results obtained with HERWIG, the other main complete implementation of BNV phenomenology. The distributions in this section are not necessarily directly observable; we will return to experimental tests in the following section.

For calculating the sparticle mass and coupling spectrum, we use the mSUGRA parameters of Snowmass points 1a and 1b [36] input to ISASUSY [2] for RGE evolution. This is done for both HERWIG and PYTHIA, and so there can be no artifacts created by non-identical EW scale superspectra when comparing the output of the two programs. Apart from this aspect, most of the topics we study here are not sensitive to the details of the SUSY parameter set

used.

For all studies, we use the “factory presets” for both PYTHIA and HERWIG. No tuning has been performed, and no demands e.g. of identical shower cut-offs or fragmentation functions have been made. This does mean that one must exercise a slight caution before drawing too strong conclusions from the comparisons, since default PYTHIA is not always directly comparable to default HERWIG. In the following, we will comment on these aspects when necessary.

5.1 Consistency checks

The main technical simplification of the junction string fragmentation scheme is the asymmetric fashion in which two of the three strings are first fragmented inwards, with the remnants joined to an effective diquark jet that is fragmented together with the third string. In particular, the fear might be that the joined diquark could produce a spurious new jet intermediate to the directions of the two strings from which it inherits the leftover momenta. Therefore the fragmentation of a symmetric three-parton configuration is studied in Fig. 9. For the production of hadrons other than the junction baryon, an essentially perfect symmetry between the three jets is obtained, both in angular and in momentum variables. As could be expected, the situation is less perfect for the junction baryon itself. Indeed, it has a tendency to be produced in a direction intermediate to the first two jets considered, rather than along either of these two separate directions, while baryons in the third jet line up the way one would like. It should be noted that most of this effect occurs for baryons of intermediate momenta $|\mathbf{p}| \simeq 1$ GeV; at lower momenta everything is and should be close to isotropic, while those few baryons that have large momenta tend to be better aligned with the jet directions. The overall balance between the jets is good, with 32% of the junction baryons found in the angular range around the third jet, relative to 34% in each of the other two. The momentum spectra in the three jets also show reasonable agreement: although the third tends to produce somewhat harder junction baryons than the first two, the difference only corresponds to a mean momentum of 0.97 GeV rather than of 0.87 GeV.

We may therefore conclude that the asymmetric algorithm does not seem to induce any imperfections in the bulk of the hadron production, and only rather modest ones for the junction baryon itself. These latter imperfections should also largely average out in realistic simulations, when the order in which jets have been hadronized is not known. Specifically, it appears to be a stable and reliable prediction of the model that the junction baryon should be found at low momentum, $|\mathbf{p}| \lesssim 2$ GeV, in the junction rest frame.

Further technical details, concerning the junction motion (subsection 4.3) and hadronization (subsection 4.4), could have been addressed in different approximations. Specifically, we have investigated how the fragmentation spectra would be affected if we made the following changes to our model:

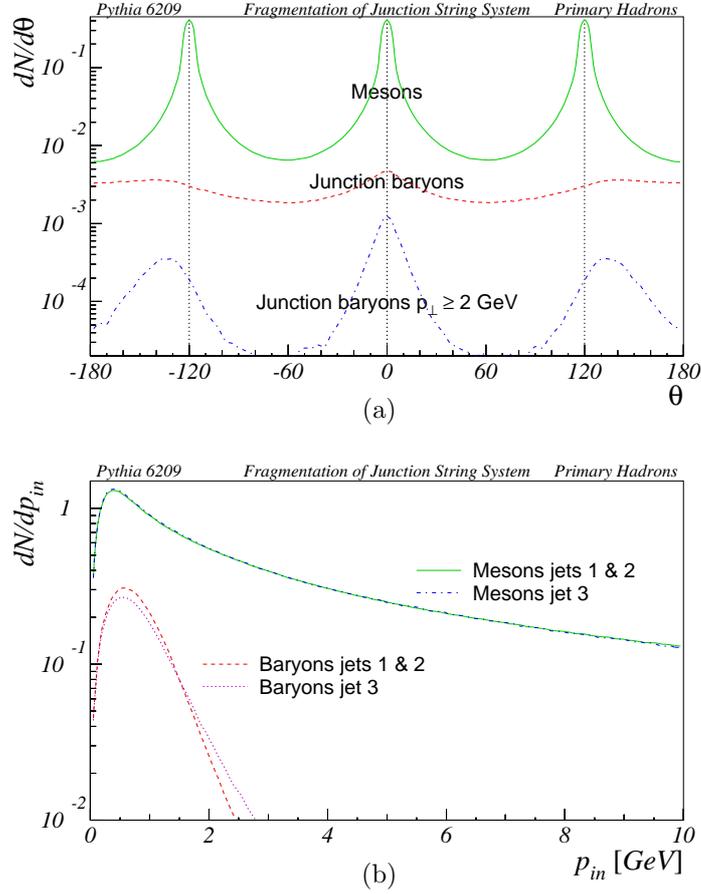


Figure 9: Fragmentation of a uuu parton configuration with 40 GeV in each jet and opening angles 120° between partons. Only the primary hadrons are shown, i.e. subsequent decays are disregarded. For simplicity all normal baryon–antibaryon production is switched off, as is the popcorn baryon possibility. (a) The particle flow $dN/d\theta$ in the event plane, for mesons, for the junction baryons and for the fraction of junction baryons with momentum in the event plane above 2 GeV. The jet at around 0° is the last one to be considered in the hadronization process. (b) The meson and junction baryon momentum spectra per jet, shown separately for the first two jets considered and for the last one. Hadron assignment to jets is based on a simple division of the plane into three equal 120° sectors.

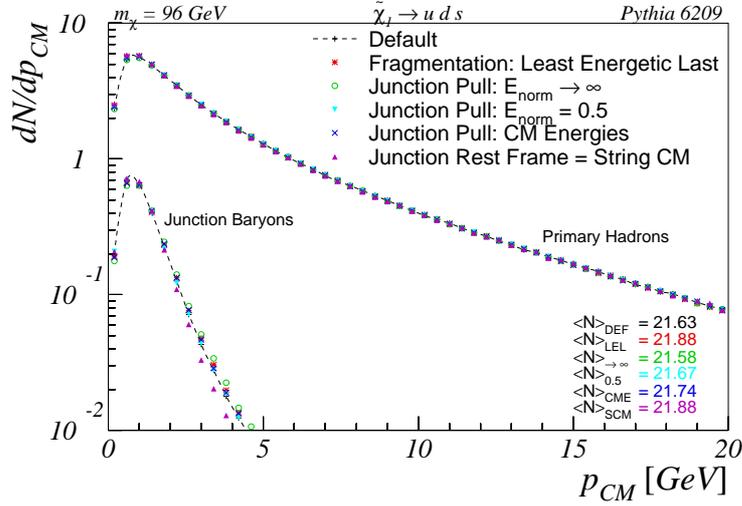


Figure 10: Momentum spectra of primary hadrons in the decays of a 96 GeV neutralino. Results with the default implementation are compared with five alternative ones. Average multiplicities of primary hadrons are shown in the lower right corner of the plot.

1. Switched the order of fragmentation of the three string pieces, so that the *last* string to be fragmented would be the one containing *least* energy, the reverse of the default behaviour.
2. Replaced eq. (13), defining the junction pull vectors used in determining the junction rest frame, by a linear sum of 4-momenta instead of the default exponentially weighted sum, equivalent to setting $E_{\text{norm}} \rightarrow \infty$ in eq. (13).
3. Enhanced the significance of soft gluon emission by setting $E_{\text{norm}} = 0.5$ GeV rather than the default $E_{\text{norm}} = 1.5$ GeV.
4. Switched off the iterative procedure for finding the junction rest frame, so that the energies appearing in eq. (13) are the string system CM energies rather than the energies in the junction rest frame.
5. Performed the fragmentation in the string system CM rather than in the junction rest frame.

The impact of these variations on the hadronization spectrum of a 96 GeV neutralino decaying to light quarks is shown in Fig. 10, where the rates of primary hadrons and junction baryons produced in the fragmentation are plotted as functions of their momenta in the CM of the decaying neutralino. Average hadronic multiplicities are shown at the lower right of the plot. As is readily observed, none of these variations lead to a significant change in the spectra;

hence we believe the systematic uncertainties associated with these assumptions to be negligible to a good precision.

As a final variation, in Fig. 10 we also study what happens if the junction is assumed to be at rest in the string system CM frame (which agrees with the neutralino CM, unless a $g \rightarrow q\bar{q}$ splitting occurred in the shower). This is not intended as a realistic model variation, but indicates that even such extreme scenarios would not change the qualitative picture of a low-momentum baryon. Here, however, we are helped by the fact that the three jets in the neutralino decay tend to be rather well separated in angle and that therefore the boost between the junction rest frame and the neutralino CM frame is not so large. Nevertheless, the more correct description gives a somewhat harder junction baryon spectrum, as should be expected from the boost.

5.2 Shower and hadronization activity

In section 3 we described the strategy for the generation of additional gluon radiation in BNV decays. Thereby the number of reconstructed jets can exceed the primary parton multiplicity. For our studies here we rely on the Durham jet algorithm [37], with a distance measure $y_{ij} = 2 \min(E_i^2, E_j^2) (1 - \cos \theta_{ij}) / E_{\text{vis}}^2$ between two particles or clusters i and j . Results can depend on the algorithm used, but we expect most of the phenomenology to come out similarly also for other algorithms [38].

In Fig. 11, jet rates are shown as functions of the y_D jet resolution parameter, with the jet clustering being performed in the center of mass of the decaying neutralino. The plots show a comparison between neutralino decay to light quarks (full lines) and neutralino decay to heavy quarks (dashed lines) for PYTHIA, at three stages of the event generation; (a) initial decay, (b) after parton shower, and (c) after hadronization.

In the initial decay, phase space is isotropically populated, and so no significant mass differences should be expected. However, since y_{ij} contains energies rather than e.g. momenta in the numerator, and since massive partons are assured nonvanishing energies, there are slightly more 3-jets for decays to heavy quarks. When the parton showers are included (cf. Fig. 11b), the suppression of gluon radiation off massive quarks ensures that decays into light quarks gives the larger number of further jets. Finally, hadronization (cf. Fig. 11c) once again flips the picture, at small resolution scales y_D , mainly by the larger p_\perp kick impaired to hadrons in the B meson decays. Such gluon radiation and hadronization effects are familiar from studies e.g. at LEP [15].

In comparing with the HERWIG implementation, we now concentrate on decays to light quarks so as to obviate the further complication of extra jets from heavy quark decays. This means that there are three main stages we can check for differences: 1) the population of the initial 3-body phase space, 2) the parton showers, and 3) the hadronization.

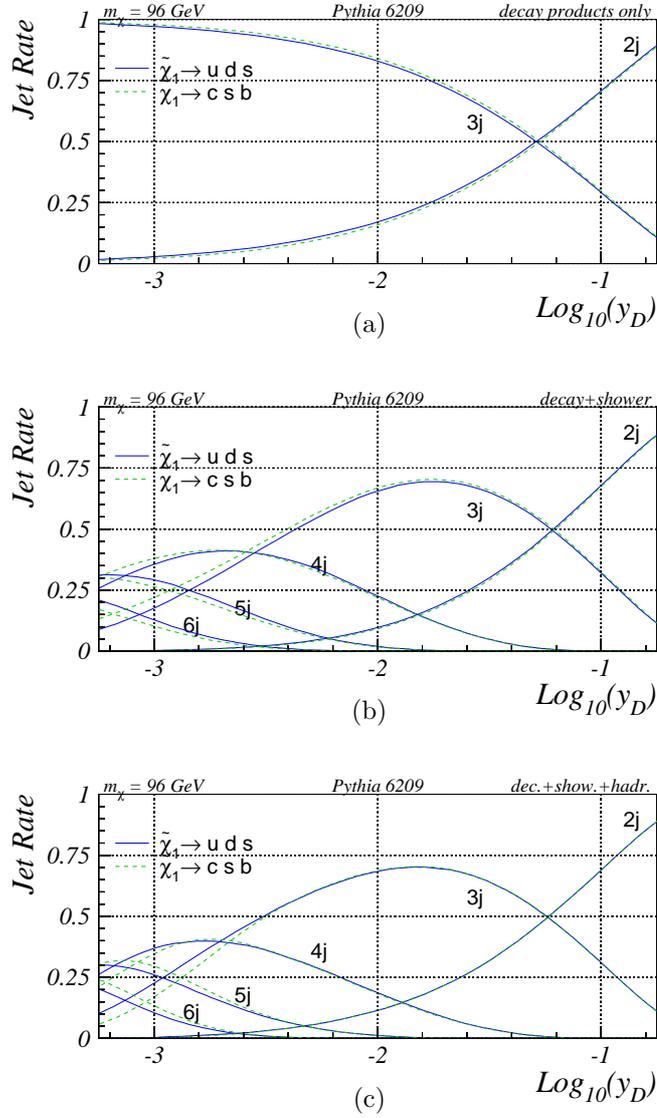


Figure 11: Jet rates in the decay of a 96 GeV neutralino resonance, $\tilde{\chi}^0 \rightarrow u_i d_j d_k$, as functions of y_D . $\tilde{\chi}^0 \rightarrow uds$ is shown with full lines and $\tilde{\chi}^0 \rightarrow csb$ with dashed. (a) Clustering only of the primary neutralino decay products, (b) clustering after showering but before hadronization, and (c) clustering after full event generation.

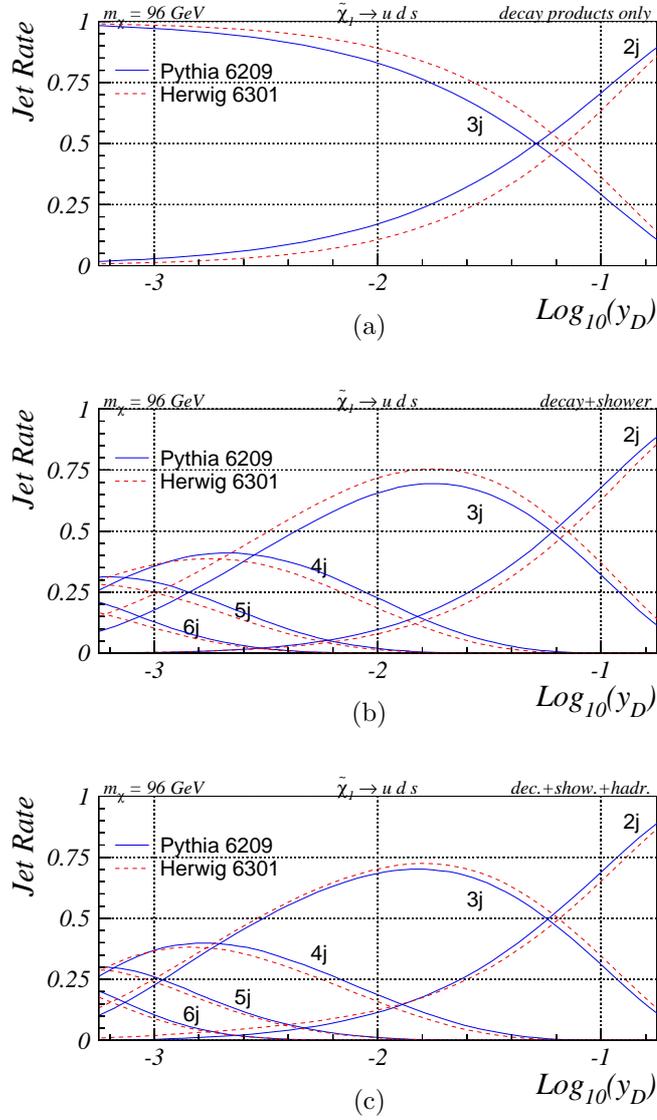


Figure 12: Jet rates in the decay of a 96 GeV neutralino resonance, $\tilde{\chi}^0 \rightarrow uds$, as functions of y_D , as obtained with PYTHIA (full lines) and HERWIG (dashed lines). (a) Clustering only of the primary decay products, (b) clustering after showering but before hadronization, and (c) clustering of the final event.

1. Phase space population. As mentioned above, PYTHIA produces an isotropic phase space population whereas HERWIG uses full matrix element weighting. In Fig. 12a a comparison of the jet rates between PYTHIA (full lines) and HERWIG (dashed lines) is shown for neutralino decay to light quarks. Only the three primary decay products are used in the jet clustering here. One notes that HERWIG has more 3-jets than PYTHIA over the whole y_D range. This agrees with expectations that the matrix elements should disfavour 2-jet configurations: in the 2-jet limit one invariant mass vanishes, meaning that contributions from graphs with that particular intermediate squark propagator also vanish. (A decay like $\tilde{\chi}^0 \rightarrow uds$ receives contributions from the three intermediate states $u\tilde{u}^*$, $d\tilde{d}^*$, and $s\tilde{s}^*$.) There are no squark poles in or near the phase space, however, so the variation of the matrix elements is rather mild. Thus effects of the matrix-element weighting are significant, but not dramatic.
2. Showering. In Fig. 12b we show the jet rates after showering but before hadronization. (For HERWIG this also means before the nonperturbative $g \rightarrow q\bar{q}$ branchings.) The difference produced in the initial decay is reduced slightly, but otherwise the pattern at high y_D is the same as above, with HERWIG producing more jets than PYTHIA. However, when the four-or-more jet rates come into play at lower values of y_D , PYTHIA tends to produce more jets than HERWIG, even if there is still qualitative agreement. We explore this difference by considering the number of partons emitted in the respective shower, Fig. 13. The variation from two- to three-jetlike topologies is characterized by the minimum opening angle between any two of the three original partons. Differences in the shower algorithms, such as the choice of cutoff scale, imply that the number of partons as such is not so informative. Therefore a convenient reference is offered by the respective results for ordinary $Z^0 \rightarrow q\bar{q}$ events, at a comparable energy. In PYTHIA, the $\tilde{\chi}$ decays produce almost a factor 3/2 more radiated partons than does the Z^0 ones. This is consistent with a rather constant rate per radiating parton, somewhat reduced by the lower energy per parton in the $\tilde{\chi}$ decays. Such a result is not unreasonable, given that most gluons are emitted in the collinear regions around the quark directions, according to a universal radiation pattern, i.e. independently of the colour flow to other quarks at wide angles. One should note that the basic parton shower formalism of PYTHIA tends somewhat to overestimate energetic wide-angle emission in $Z^0 \rightarrow q\bar{q}$ events. This is then compensated by a rejection factor that reduces this emission to what is expected from the $q\bar{q}g$ matrix elements [18, 19]. Since no corresponding matching is implemented for $\tilde{\chi} \rightarrow qqq$ decays, it is quite likely that the shower activity here is somewhat overestimated. Further, in the limit that two partons become collinear, one should expect them partly to screen each other and the overall activity thereby to drop down towards the Z^0 level. That this does not happen in

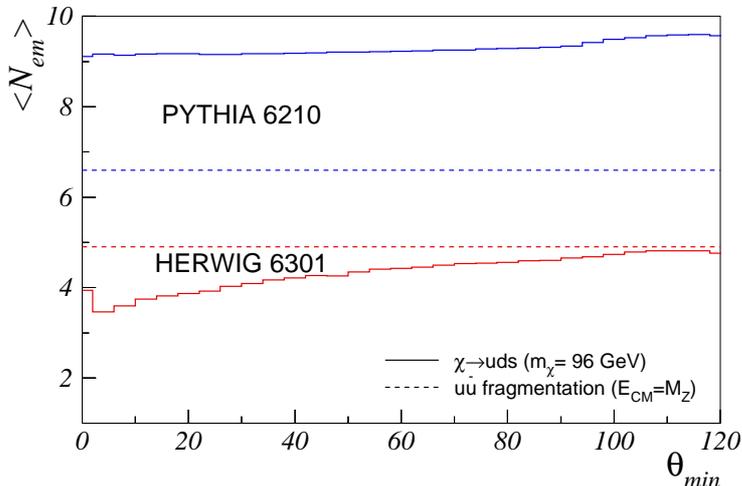


Figure 13: The average number of partons emitted in the shower of $\tilde{\chi} \rightarrow uds$ decays, as a function of the smallest opening angle between the neutralino decay daughters (solid lines; upper: PYTHIA, lower: HERWIG). The average number of partons emitted in $u\bar{u}$ cascades at the Z^0 resonance is shown for comparison (dashed lines).

PYTHIA is an obvious shortcoming of the modelling, which is not set up to handle these screening effects.

By contrast, in HERWIG the variation with event topology is more marked: the HERWIG radiation cone is defined by the opening angle to one of the other quarks picked at random [8], meaning that two nearby partons may kill the radiation from each other. The most notable feature, however, is that the $\tilde{\chi}$ multiplicity is lower than the Z^0 one, also for well separated quark topologies, contrary to the arguments above.

We do not claim to understand the low emission rate in HERWIG. One observation, that we will quantify below, is that $\tilde{\chi}$ decays tend to have a large-mass central cluster, which could indicate a shortfall of emissions at wider angles. Among several potential differences between the two shower descriptions, it is then tempting to suspect the procedures used to start up the showers from some maximum virtuality. As we have already noted, the “dead zone” phenomenon is a well-known reason for HERWIG to underestimate multi-jet rates. If so, the differences should be larger for out-of-the-event-plane distributions than for inclusive ones, since the direction at 90° to the event plane is as far away as one can be from any of the three original quark directions. In Fig. 14, distributions of the linearized Aplanarity event shape variable [39] are compared. The curves relevant

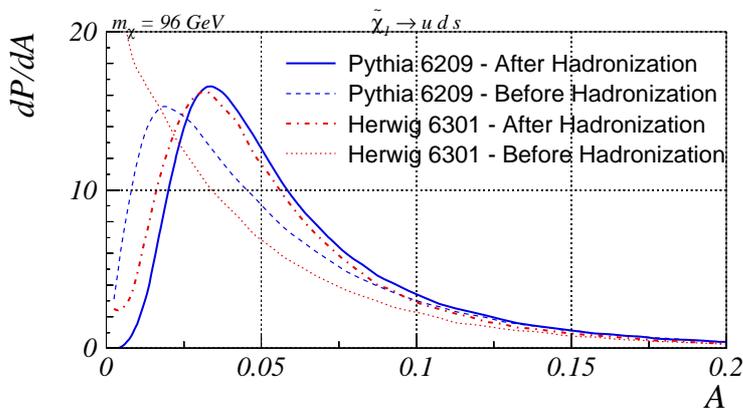


Figure 14: Distribution of linear Aplanarity, comparing PYTHIA and HERWIG.

for showering are the dashed for PYTHIA and dotted for HERWIG. One observes that, also in the present case, HERWIG generates less activity out of the event plane than does PYTHIA. The steep rise of the HERWIG curve towards small Aplanarities is caused by events where no emissions harder than roughly 1 GeV or so occurred. The pattern is consistent with the hypothesis of a shortfall of wide-angle emissions in HERWIG.

3. Hadronization. The results of jet clustering after full event generation are shown in Fig. 12c. One notes that HERWIG still produces more jets at high y_D , whereas PYTHIA has more jets at small y_D , but overall the agreement between the two programs here is impressively and unexpectedly good. In particular, the hadron-level jet rate in HERWIG goes up appreciably relative to the parton-level one, whereas there is less difference between the two in PYTHIA.

Of special relevance here is the HERWIG treatment of the “baryon cluster” which carries the net baryon number generated by the BNV decay. As we see from Fig. 15, this cluster has a rather different mass spectrum than other clusters, with a much higher average mass. (This is also noted in [8].) Such a large-mass cluster is first fragmented into smaller clusters, which each then decay to hadrons. The fragmentation is along an assumed “string” direction, as appropriate for simple $q\bar{q}$ clusters. In the current case, where three leftover quarks are to be considered, two of these, picked at random, are combined to a diquark. The resulting system is fragmented along a single quark–diquark axis. The extreme case is those events, roughly 2% of the total number, where the whole neutralino becomes a single cluster. Then, as shown in Fig. 16, no well-separated third jet exists after hadronization, this information having been erased by the

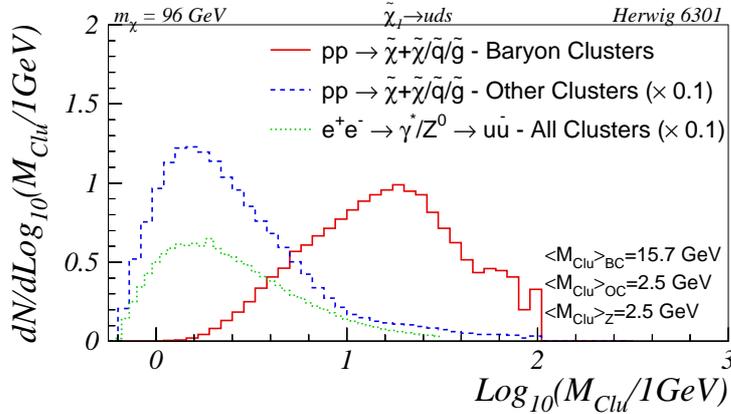


Figure 15: HERWIG cluster mass spectra for the clusters containing non-zero baryon number (solid lines) as compared to the ordinary clusters from the neutralinos and the additional perturbative activity in the event. For reference, the cluster spectrum for hadronic Z^0 decays at 90 GeV are also shown (dotted lines). Only the baryon cluster spectrum is shown with its true normalization, the two other spectra having been normalized by a factor 1/10 to fit them onto the same scale as the baryon clusters.

cluster formation procedure! In a normal event, effects tend in the opposite direction: some clusters appear along each of the three quark directions as a consequence of the shower activity, but then the baryon cluster also emits particles in a new “jet” direction intermediate to the two quarks that get combined into a diquark. The net result is an increase in the non-perturbative jet rate relative to the perturbative one. To the extent that this effect is unintentional, the large hadronization correction in HERWIG is misleading.

To summarize, we note that the final agreement in jet rates must largely be viewed as coincidental, since all three contributing physics components by themselves differ between HERWIG and PYTHIA.

5.3 The junction baryon

As discussed in subsection 4.4, the non-perturbative collapse of the colour wavefunction, which in our model results in a Y-shaped string topology, essentially “traps” the baryon number around the locus of the string junction. Thus, in the rest frame of the junction, we expect junction baryon momenta of the order of the hadronization scale. For well-separated jets, we further expect the junction rest frame to be only slightly different from the CM frame of the decay. Thus a generic feature of our model is the prediction that the baryon number

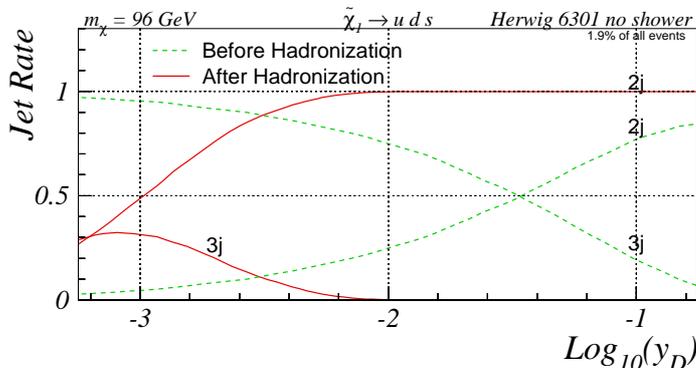


Figure 16: 2- and 3-jet rates before and after hadronization for HERWIG events with no gluon emission ($\sim 2\%$ of the event rate for default HERWIG).

generated by BNV decays predominantly ends up in baryons which have small momenta relative to the decaying particle.

In Fig. 17 the rates of both junction baryons and “ordinary” baryons are plotted as functions of their CM momenta for $\tilde{\chi} \rightarrow uds$ for (a) two different neutralino masses (Snowmass points 1a and 1b), and (b) comparison between PYTHIA and HERWIG for a 96 GeV neutralino.

As expected, one observes a much sharper peaking towards low momenta for the junction baryon distributions as compared to the ordinary baryon distributions in Fig. 17a. Also, note that this peaking does not depend on whether the neutralino is light or heavy. A heavier neutralino gives rise to longer strings, i.e. larger momenta and higher multiplicities, as can be seen by the shift in shape and normalization of the ordinary baryon momentum distribution, but this does not significantly affect how much energy will eventually be available to form the junction baryon. The slight hardening of the junction baryon momentum for a higher neutralino mass is attributable to the increased phase space for perturbative gluon emission, which gives the junction a slightly higher average velocity in the neutralino rest frame.

In Fig. 17b we compare with the HERWIG implementation. Of course, HERWIG does not have junctions and so no junction baryons *per se*, but it is nonetheless possible in most cases to trace the BNV-associated baryon number, via a cluster with non-vanishing baryon number, to a specific final-state baryon. And so, with a slight abuse of nomenclature, we term this baryon the HERWIG “junction baryon”. Naturally, both programs produce only one junction baryon per decay. However, as shown in Fig. 17b, the HERWIG fragmentation produces 1.1 ordinary baryons more, on the average, than does PYTHIA. Interpreted as a systematic uncertainty on the baryon multiplicity in fragmentation, this difference is disturbingly large. However, when comparing with the LEP ex-

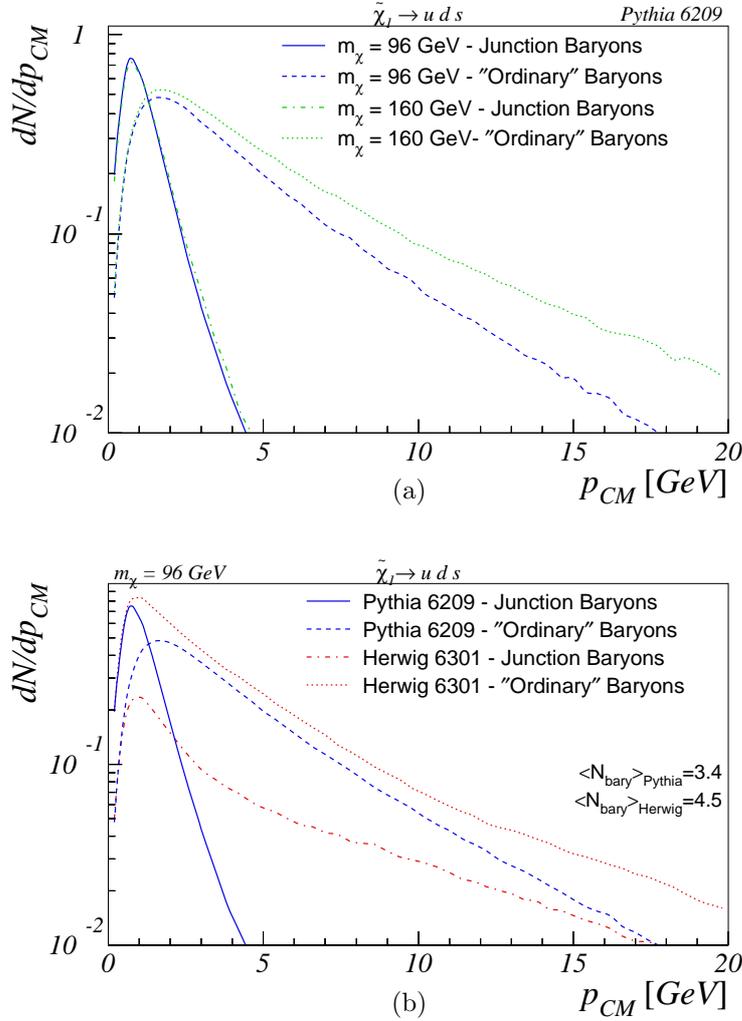


Figure 17: CM momentum distributions in the decay of a 96 GeV neutralino for junction (full line) and ordinary (dashed line) baryons (a) as compared to the decay of a heavier neutralino (dot dashed and dotted line for junction and ordinary baryons respectively), and (b) as compared to the same process in HERWIG (dot dashed and dotted line for junction and ordinary baryons respectively).

perimental value at the Z^0 peak, $\langle n_p \rangle_{\text{LEP}} = 0.98 \pm 0.1$ [15] protons per event (neutrons not being measured), default PYTHIA with its 1.2 protons per event

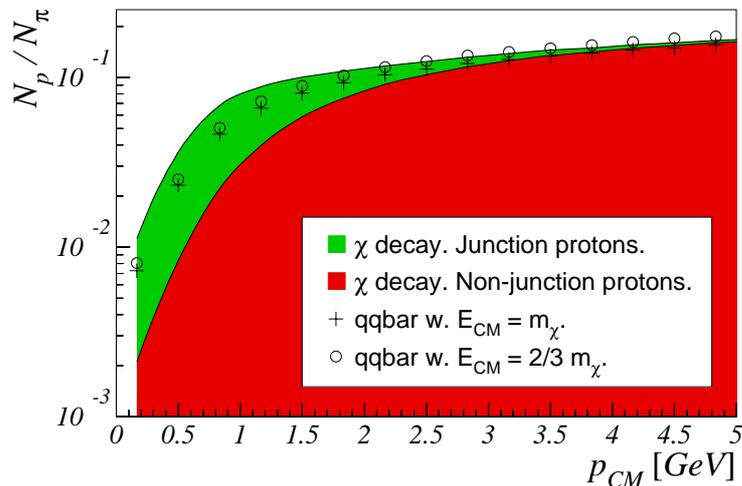


Figure 18: Proton (+antiproton) CM-frame momentum spectrum normalized to the charged pion one. Distributions are shown for 96 GeV neutralino decays to light quarks (junction and non-junction proton contributions shown separately) and for two cases of ordinary $q\bar{q}$ fragmentation, one (crosses) where the total CM energy is the neutralino mass and one (circles) where it is only $\frac{2}{3}m_{\tilde{\chi}}$, in order to make the energies of each piece of string equal.

comes closer than default HERWIG with its 1.5. Presumably a better tuning to LEP than offered by the defaults would also reduce the difference seen in neutralino decays.

The difference in the description of the junction baryon should persist, however. In HERWIG there is no particular reason why this baryon should be slower than other baryons, and this results in a distribution for junction baryons which looks more or less the same as that for ordinary baryons, a situation quite at contrast to the shape of the junction baryon spectrum predicted by the model presented here, cf. Fig. 17b. This means that the HERWIG model offers less hope to find direct evidence for the occurrence of BNV. We believe, however, that this model is unrealistic, being flawed by an underestimated shower activity in combination with an unphysical description of the fragmentation of the baryon cluster, as explained above in section 5.2.

When comparing baryon rates between physical processes, differences in the overall multiplicity may confuse the issue. For instance, a $\tilde{\chi} \rightarrow qqq$ decay should have a larger multiplicity and softer fragmentation spectra than a $\gamma^*/Z^* \rightarrow q\bar{q}$ decay of the same mass, since the total energy is shared between

more strings. A decay to two gluons, such as $h^0 \rightarrow gg$, would have an even larger multiplicity, owing to the gluon energy being shared between two strings and, by the same colour charge argument, also radiating more in the perturbative phase. A more realistic measure is the fraction of (anti)protons among charged particles, where the issue of overall multiplicity divides out, and quark and gluon jets of different energies give almost identical results at low momenta. In Fig. 18 we show the ratio of proton to pion momentum spectra, which is basically equivalent with the above. (Inclusion of kaons along with the pions would not affect the relative difference between scenarios.) The mass differences and decay patterns give a softer pion than proton spectrum, and therefore a rising p/π ratio. In the neutralino decay, the significant enhancement caused by the junction baryon contribution is again visible. It is noteworthy that the case of a $q\bar{q}$ (or gg) decay gives a ratio intermediate to the ones in the qqq decay without and with the junction baryon included. The explanation is that the low-momentum region in the qqq topology is depleted from normal baryon production by the presence of the junction baryon. Unfortunately, this means experimental signals are not quite as dramatic as might have been wished for, as we shall see.

5.4 Generic event properties

We now turn our attention to a comparison of the expected overall event features between PYTHIA and HERWIG, specifically as characterized by the fragmentation spectra and the multiplicity distributions for charged hadrons.

In Fig. 19 we show the hadronic fragmentation spectra produced by the decay $\tilde{\chi} \rightarrow uds$. With “primary hadrons” and “final state hadrons” we understand hadrons produced in the fragmentation of the strings, before and after decays take place, respectively. There is good agreement, with slightly smaller primary hadron momenta being produced by HERWIG as compared to PYTHIA.

In Fig. 20 is compared the multiplicities of charged hadrons as produced by PYTHIA and HERWIG. An average difference of almost 2 charged hadrons can be observed. When performing the same comparison for the LEP1 process $e^+e^- \rightarrow \text{hadrons}$, where $\langle n_{ch} \rangle_{\text{LEP}} = 20.92 \pm 0.24$ [15], the PYTHIA 21.0 and HERWIG 20.1 values also differ. This mistuning partly explains the difference for neutralino decay between the two programs. One should remember that the two hadronization models for this kind of configurations are quite different from the respective description of Z^0 decay, however, so there is no strong reason for the two programs to agree, neither with each other nor with the Z^0 numbers.

Consider e.g. the simple string model, which predicts a logarithmic increase of the multiplicity with jet energy, $\langle n \rangle \propto \log(E/m_0)$, where m_0 is some typical hadronic mass scale of order 1 GeV. Taking LEP events to consist of two jets with half the energy each, and neutralino decays to be three jets with a third each, the above LEP1 multiplicity would translate into something like 28

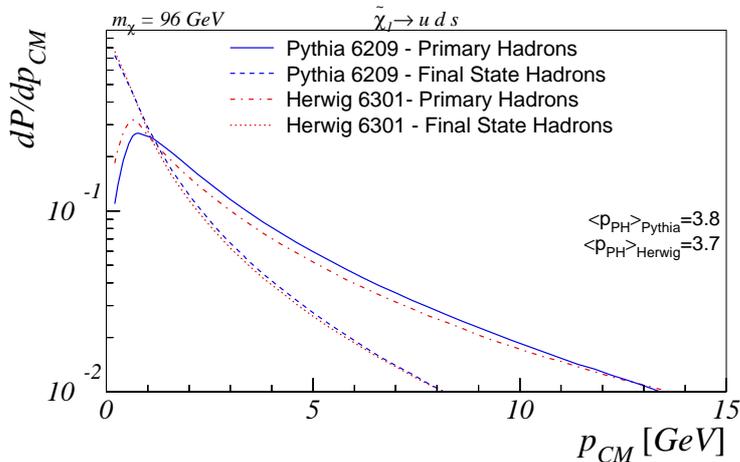


Figure 19: CM momentum distributions in the decay of a 96 GeV neutralino for primary hadrons (PYTHIA full lines, HERWIG dot-dashed) and final state hadrons (PYTHIA dashed lines, HERWIG dotted). At the bottom left, the mean primary hadron momenta are shown for each of the two programs.

charged particles per neutralino decay. That the programs predict multiplicities much below this indicates the importance of the details of the event description.

5.5 Alternative string topologies

Until now, we have considered only neutralino decay, since most of the salient features of our model stand out more clearly in this case. Nevertheless, in regions of SUSY parameter space where the BNV couplings are either larger than the gauge couplings, or where the gauge decays of the lowest-lying squark mass eigenstate are kinematically suppressed or forbidden, a special situation arises: the direct 2-body BNV decays listed in section 2 dominate for the lowest-lying squark mass eigenstate. This gives us a unique chance to study the more exotic string topologies discussed in section 4.5: junction–junction strings and junction–junction annihilation.

To quantify, we here take a generic “light stop” SUSY spectrum, Snowmass point 5, with $m_{\tilde{t}_1} = 220$ GeV, and the decay $\tilde{t}_1 \rightarrow \tilde{\chi}_1^+ b$ kinematically inaccessible. Since we wish specifically to study 2-junction and 0-junction topologies, we assume third generation BNV couplings of order 0.1 or larger, so that the stop lifetime, cf. eq. (14), is sufficiently small that string breaks between the stops will not normally occur before decay. Furthermore, we require that the stops be produced in an overall colour-singlet state, so that their decays, leaving out

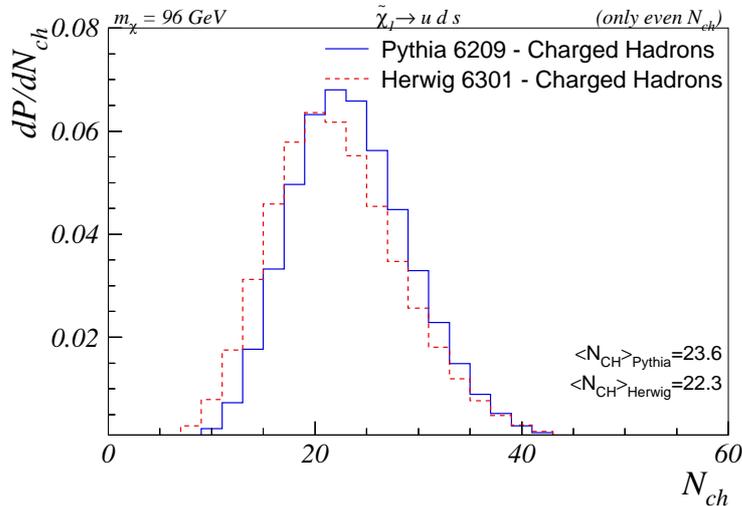


Figure 20: Charged hadron multiplicities in the decay of a 96 GeV neutralino as obtained with PYTHIA (full lines) and HERWIG (dashed). At the bottom left, the mean charged hadron multiplicities are shown for each of the two programs.

the possibility of colour reconnections in the final state, are colour-connected to each other. This will generally not be the case at hadron colliders since, in processes like $gg \rightarrow \tilde{t}_1 \tilde{t}_1^*$, each stop inherits its colour from a parton in the initial state rather than from the other stop.

Consequently, we consider the process $e^+e^- \rightarrow \tilde{t}_1 \tilde{t}_1^*$, with the stops decaying to light quarks, $\tilde{t}_1 \rightarrow \bar{d}s$. Again, differences with respect to the HERWIG implementation exist, but since we expect the main part of these to be covered already by the discussion in the preceding subsection, we do not perform explicit comparisons here. Suffice it to note that the hadronization model for BNV squark decays adopted in HERWIG roughly parallels that of PYTHIA with only 0-junction configurations.

As was observed in subsection 4.5, the probability for two connected junctions to annihilate is almost unity close to threshold in our model. To ensure that the case where we let the choice between a 0-junction and a 2-junction topology be determined dynamically, i.e. by eqs. (20)–(23), indeed does differ from the pure 0-junction case, we choose a CM energy somewhat above threshold. For Snowmass point 5, the threshold is $2m_{\tilde{t}_1} = 440 \text{ GeV}$, thus a CM energy of 800 GeV is appropriate. This results in a junction–junction annihilation rate of about 0.6, cf. Fig. 8.

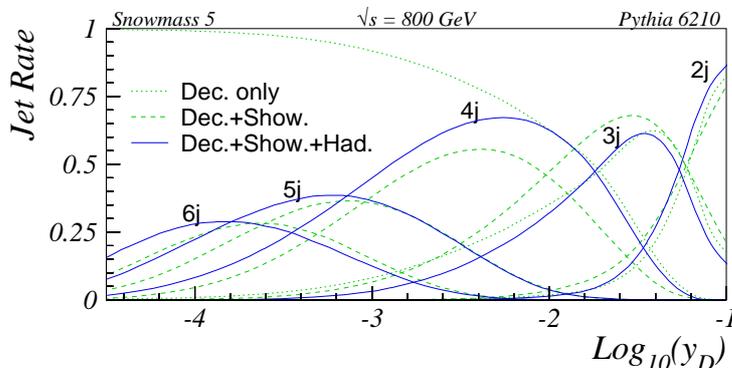


Figure 21: Jet rates for the process $e^+e^- \rightarrow \tilde{t}_1(\rightarrow \bar{d}\bar{s})\tilde{t}_1^*(\rightarrow ds)$ at 800 GeV CM energy, as functions of y_D , for three stages of the event generation: clustering only of the primary $\tilde{t}_1\tilde{t}_1^*$ decay products (dotted curves), clustering after showering but before hadronization (dashed curves), and clustering after full event generation (solid curves).

We again begin by considering the jet rates, for the specific process $e^+e^- \rightarrow \tilde{t}_1(\rightarrow \bar{d}\bar{s})\tilde{t}_1^*(\rightarrow ds)$ at 800 GeV CM energy. The jet clustering is performed using a Durham algorithm in the e^+e^- CM frame, yet since we here make the simplification of neglecting initial-state radiation and beamstrahlung, this frame is identical to the $\tilde{t}_1\tilde{t}_1^*$ CM. Note that this frame need not generally coincide with the $\tilde{t}_1\tilde{t}_1^*$ CM at decay time, due to the possibility of the stop pair emitting final-state radiation.

In Fig. 21 we show the jet rates for dynamical selection between 2-junction and 0-junction topologies, as functions of y_D . To trace the evolution of the event, we plot the rates for initial decay (dotted curves), after parton shower (dashed curves), and after hadronization (solid curves).

Since each stop decays to a 2-body final state, we obtain four jets at leading order. As can be observed from the dotted curves in Fig. 21, the initial decay daughters are actually clustered to fewer than 4 jets for a non-negligible fraction of the events down to quite low values of y_D , this due to asymmetric decays causing two daughters to end up close to each other and so be clustered to one jet rather than two.

The question now arises to what extent, if any, the choice of string topology affects the jet rates. To study this, in Fig. 22 we replot the solid curves of Fig. 21 together with the two extreme cases where all events are forced to be either of the 2-junction (dotted curves) or the 0-junction (dashed) type. Note that the latter sample has a small contamination of 2-junction events, from perturbative breakups in the $\tilde{t}_1\tilde{t}_1^*$ shower. As can readily be observed, the curves are identical to a high precision, bearing witness to the absence of any resolvable structural

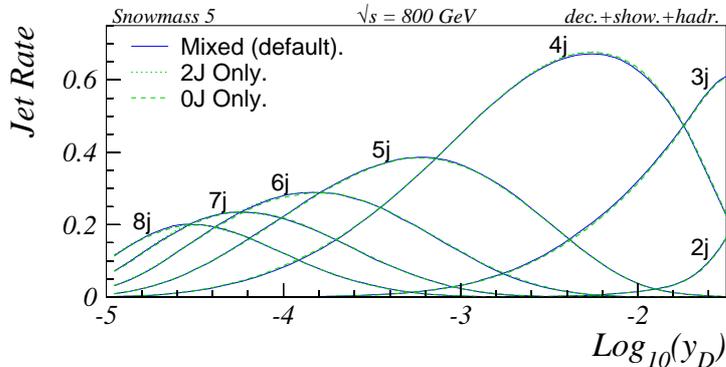


Figure 22: Jet rates for the process $e^+e^- \rightarrow \tilde{t}_1(\rightarrow \bar{d}\tilde{s})\tilde{t}_1^*(\rightarrow ds)$ at 800 GeV CM energy, as functions of y_D , for three choices of string topology: mixed 0-junction and 2-junction (solid curves), 2-junction only (dotted curves), and 0-junction only (dashed curves).

difference between the topologies, at least for the y_D ranges and jet numbers we found it meaningful to include in the analysis.

We thus do not believe that these semi-inclusive observables could be used e.g. to measure the junction–junction annihilation rate. We return to the issue of whether exclusive quantities can do a better job in the more experimentally oriented study in subsection 6.3.

6 Experimental tests

In this section we present some semi-realistic studies on how baryon number violation could be pinned down experimentally. We do not propose to recapitulate the extensive experience on SUSY search strategies, with likelihood techniques or a succession of cuts and sideband subtraction procedures, required to isolate signals relative to backgrounds for different SUSY parameter sets. Rather, we assume that a signal has already been found in multijet final states, with bumps in the two- or three-jet invariant mass distributions consistent with the presence of new particles. The key issue is then whether this is evidence for SUSY with BNV or for some completely other physics signal. The decision is likely to be based on a combination of different pieces of information, but a crucial one would be to find the baryon associated with the BNV itself. This is not possible on an event-by-event basis, owing to the background of normal baryon production and the non-identification of neutrons. What we can aim for is then to demonstrate the existence of a statistically significant excess (or not) of baryon production in the momentum region predicted by our

model.

This kind of demonstration will most likely be beyond the reach of the Tevatron; therefore in the following we concentrate on an LHC scenario with pp collisions at $\sqrt{s} = 14$ TeV. We take a crude representation of the ATLAS tracking and calorimetry capabilities (described below) to be sufficiently detailed for the kind of analysis we will consider, specifically whether an excess of protons à la Fig. 18 can still be seen for fully generated events in a semi-realistic detector environment.

Of course, a similar exercise could have been performed for CMS but, at our level of detail, we expect the differences between the two experiments to be of negligible consequence. The exception would be baryon identification strategies, see below.

With respect to the prospects at linear colliders, the comparatively clean environment and high luminosity imply that studies similar to the ones presented here would be easier to carry out, although they would obviously have to wait for such a facility to be built. Therefore, we constrain our attention in this direction to one special case, that of colour-connected BNV $\tilde{t}_1\tilde{t}_1^*$ decays, which we expect to be nigh impossible to study with hadron colliders.

6.1 Detector layout and acceptance

Based primarily on the ATLAS physics TDR [40], we assume a calorimeter that covers the region $|\eta| < 5$ in pseudorapidity and $0 < \phi < 2\pi$ in azimuth, with a granularity of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$. This reproduces the expected ATLAS calorimeter granularity at mid-rapidity, but overestimates it at large rapidities. Smearing due to finite resolution effects has not been simulated.

Proton and pion reconstruction. For p^\pm and π^\pm reconstruction and identification we use the region $|\eta| < 2$, well inside the coverage of the ATLAS and CMS tracker systems. Inside this region, all charged hadrons with $p_\perp > 1$ GeV are reconstructed. This is somewhat more optimistic than what would be possible in the real world. Particles with $p_\perp < 1$ GeV predominantly come from the underlying event, so the p_\perp cut reduces the study to particles associated with the hard physics of the event.

Ideally, we would like to identify all charged hadrons in the above region, with complete separation between the p^\pm signal and the π^\pm and K^\pm background. In reality this will never be possible. Neither ATLAS nor CMS has yet carried through detailed detector simulation of their proton identification performance. Our physics scenario may represent the first example where this capability is relevant in the study of physics beyond the Standard Model, and so could offer a convenient target for future detector studies. (QCD physics obviously provides several reasons to have this capability, one of which will be mentioned later.) In the meantime we exemplify what could be possible.

The conservative approach is to restrict the search to the $\Lambda^0 \rightarrow p^+\pi^-$ channel which, due to its two charged tracks from a displaced vertex and narrow mass peak, offers a reasonably clean signal. This channel also has the advantage that it doesn't cut down the statistics much: the flavour composition of fragmentation in e^+e^- annihilation is measured to give $N_{\Lambda^0}/N_p \sim 0.4$ [41] and the branching ratio of Λ^0 to $p^+\pi^-$ is about 60% [41], so about one in every four protons comes from a Λ^0 decay to $p^+\pi^-$. Since the proton inherits most of the momentum in the Λ^0 decay, the proton plots shown below would also apply for Λ^0 , with minor modifications.

Nevertheless, one may aggressively assume some proton identification capability, to augment the statistics from the Λ^0 decay, and maybe also to clean up this signal from the $K_S^0 \rightarrow \pi^+\pi^-$ background. For ATLAS, charged hadron identification is not a prime objective for the Inner Detector, but the ionization loss, dE/dx , in the Transition Radiation Tracker gives enough information to separate protons from pions by more than 1σ in the momentum range $3 \text{ GeV} \lesssim p \lesssim 20 \text{ GeV}$ [40]. However, the K^\pm/p^\pm separation is very poor even in this range, so an analysis purely based on this capability will contain a non-vanishing misidentification background. Calorimetry provides auxiliary information, especially for antiprotons, which release an additional GeV of energy by their annihilation. This should be noticeable, at least for antiprotons stopping in the electromagnetic calorimeter. For protons the mismatch between energy and momentum goes in the other direction and is smaller. We note that the hadrons discussed here are not ones found in the core of jets, and anyway have so low momenta that they are significantly deflected by the magnetic field. Therefore they are likely to be isolated in the calorimeter.

The possibility to use dE/dx in the CMS silicon tracker is still under study, but in any case the tracker has only a few sampling points, and so limited capability in this respect. On the other hand, there is considerable empty space between the silicon layers. Therefore one can imagine to install a time-of-flight system at some future date. This would allow good proton identification in the lower part of the interesting momentum range. As above, calorimetry could be used to identify some of the antiprotons.

Jet reconstruction. We use a cone algorithm over the full fiducial volume of the calorimeter, $|\eta| < 5$, with a cone size of $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} = 0.4$ and a minimum jet seed energy of 5 GeV. Reconstructed jets with a transverse energy below 25 GeV are not used in the further analysis. No jet energy recalibration procedure has been applied.

6.2 Physics cuts

As mentioned above, we assume that a relatively clean sample of SUSY events has been isolated before the present analysis is applied, hence we do not study

SM backgrounds. The event sample thus consists of a general SUSY simulation for our reference point, Snowmass point 1a [36], including all SUSY processes currently implemented in PYTHIA. Most of these, however, are suppressed at hadron colliders so that the cross section is dominated by gluino and squark production. We concentrate on the case where the BNV couplings are smaller than the gauge couplings, thus sparticles will only use the BNV channels as a “last resort”. That is, the ordinary MSSM pattern of cascade decays will persist, with the modification that the LSP decays. Assuming that the lightest neutralino is the LSP, we are thus looking at events of the type:

$$2(q/g) \rightarrow 2(\tilde{q}'/\tilde{g}) \rightarrow 2(q''/g)(+X) + 2(\tilde{\chi}_1^0) \rightarrow 2(q''/g)(+X) + 2(q_i q_j q_k), \quad (24)$$

where the two initial partons come from the beam protons and X signifies what, if anything, is split off in the cascades besides the neutralinos. Since the neutralinos are not colour connected to the rest of the event, we note that this type of event is nothing but an ordinary R -conserving MSSM cascade event with the missing energy transformed into two separate $\tilde{\chi}_1^0$ decays.

For such events, about 25% of the junction protons lies within the detector acceptance. Without any further cuts, the equivalent of Fig. 18 looks as depicted in Fig. 23. A rather depressing result, of course due to the many protons which originate exterior to the neutralino decay itself, and to the fact that we are looking at momenta in the CM of the original collision rather than in the CM of the neutralino.

However, it is possible to do much better by placing cuts designed to identify the 3 jets from each of the $\tilde{\chi}_1^0$ decays, reconstructing the presumed $\tilde{\chi}_1^0$ CM and looking for low-momentum protons in that frame. In principle, it would be even better to attempt to reconstruct the junction rest frame itself, based on the directions of the three jets. Here, we note that for well separated jets, as seen in the neutralino CM, the junction rest frame is at any rate closely approximated by the CM frame, and so in this preliminary study we content ourselves with working in the 3-jet CM.

We use two different strategies to identify the neutralino decay jets. The first, blind 3-jet reconstruction, probably does worse than what could be done with more elaborately designed cuts, and the second, optimized 3-jet reconstruction, probably errs on the side of too optimistic results. Thus, the actual experimental result should lie somewhere between these two extremes.

Blind 3-jet reconstruction

We term this analysis “blind” since we assume no *a priori* “divine knowledge” of which jets are the correct ones etc. Thus, we use a selection of cuts to attempt to identify the neutralino decay jets as precisely as possible, then reconstruct the presumed neutralino CM and look for low-momentum particles there. Since our cuts have not been through a long, rigorous optimization procedure, it is

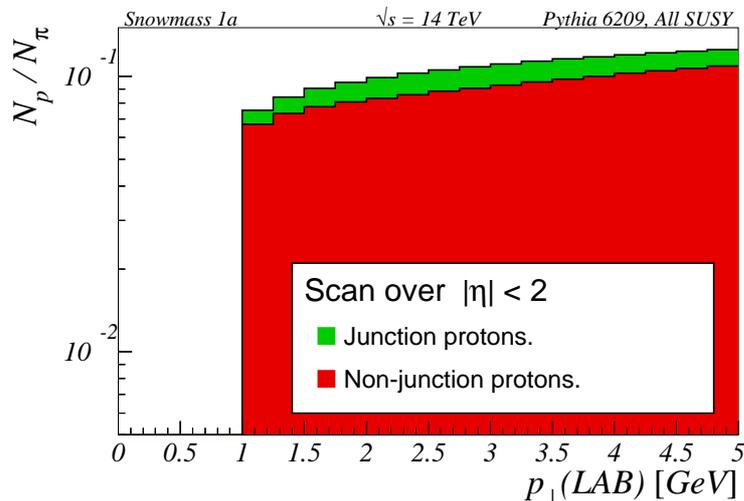


Figure 23: Reconstructed p_{\perp} spectrum for the ratio of protons to charged pions in a simple scan over the detector volume.

likely that we err on the low side of what could be done with better optimized cuts.

Note also that our cuts rely to some extent on features specific to the SUSY scenario we use (Snowmass point 1a), but they should produce acceptable results for any scenario that shares its general features: a neutralino LSP and a rather light sparticle spectrum with a relatively large gap between the squarks (and the gluino) and the LSP.

With this type of spectrum, the squark/gluino decays release much kinetic energy, giving rise to generally high- p_{\perp} neutralinos. Since the $\tilde{\chi}_1^0$ mass itself is small in comparison, its three decay jets always occupy a relatively small total region of the detector. These two considerations form the backbone of our cuts. In decays like $\tilde{\chi}_1^0 \rightarrow csb$, one could furthermore require heavy-quark tagged jets, but we do not rely on such extra information here.

To quantify, we select events with 8 or more reconstructed jets and in those look for systems of 3 jets which have a total $p_{\perp}^{3j} > 200$ GeV. The measure we use for how far the three jets are from each other in the detector is the maximal jet–jet R distance:

$$\Delta R_{ij}^{\max} = \text{Max}(\Delta R_{j_1 j_2}, \Delta R_{j_2 j_3}, \Delta R_{j_3 j_1}) \quad (25)$$

where j_i stands for the i 'th jet in the candidate 3-jet system. We impose a

cut requiring $\Delta R_{jj}^{\max} < 0.8$. Since we use a cone size of $\Delta R = 0.4$, this cut translates into looking for 3 jets which overlap or just touch each other. The details of how the energy in the overlapping regions is assigned to the each of the overlapping jets depends on the particular jet algorithm used. However, as long as double counting is avoided, we do not expect these details to influence our analysis; ideally, the CM of the 3 jets should be invariant to changes in the sharing of energy between its 3 constituent jets.

To further bring down the combinatorics, we require that none of the 2-jet pairs in the 3-jet configuration have an invariant mass close to or larger than the neutralino mass. Assuming that this is already known to some precision, we reject candidate 3-jet systems where any pair of jets has a mass larger than 90 GeV. Of the remaining candidate 3-jet systems, we select those which reconstruct to the neutralino mass ± 10 GeV.

For each of these 3-jet systems we search for candidate junction protons within the detector acceptance described above. We check whether there are any (anti)protons within $\Delta R_{p3j} < 0.5$ of the 3-jet system momentum direction. Any such protons are accepted as candidate junction protons. About 5% of the junction protons within the detector acceptance survive these cuts. The corresponding number for the non-junction protons is 0.2%.

The interesting quantity now is the momenta of these protons in the 3-jet system CM frame, our first approximation of the junction rest frame. As has been argued above, we expect that the real junction protons will exhibit a softer momentum spectrum in this frame than will be the case for non-junction protons. Thus, we expect that the enrichment of the sample by junction protons should be largest for small proton momenta in the reconstructed neutralino CM. Results are depicted in Fig. 24.

The situation has improved noticeably over the one in Fig. 23, but there is still a lot of non-junction protons at small p_{\perp} which, as we shall see in the next subsection, may be further reduced by reducing the combinatorial background, i.e. by having a more precise identification of the correct neutralino decay jets.

Optimized 3-jet reconstruction

We now assume that a nearly perfect neutralino decay jet identification can be performed, by applying some selection of cuts, the details of which are not interesting here. Rather, we mimic the effects of such optimized cuts by using the event generator information.

For each of the neutralino daughters, we select that jet which lies closest in \vec{p}_{\perp} and use only those jets for reconstructing the presumed $\tilde{\chi}_1^0$ CM's. The combinatorics are thus brought down to very low levels. One technicality involved is how we deal with situations where one jet turns out to be the closest in \vec{p}_{\perp} to more than one of the neutralino daughters. This situation can arise for instance if those two daughters were sufficiently close together to be clustered

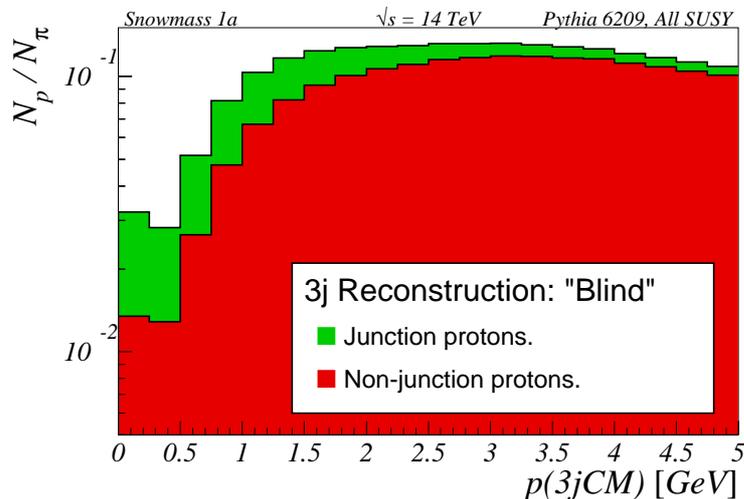


Figure 24: Reconstructed proton momentum spectrum in the 3-jet CM frame for the “blind” neutralino decay jet identification.

to a single jet, or if one of the daughters produced a jet that lies outside the active calorimetry range. Such systems are not included in the analysis.

For each accepted 3-jet system we search for candidate junction protons and charged pions in exactly the same manner as for the blind analysis. Only 1% of the junction protons within the detector acceptance survive, still with a contamination of about 0.2% of the non-junction protons, but the junction proton distribution is now more sharply peaked towards low momenta (in the 3-jet CM frame) and the non-junction protons are found at higher momenta.

The distribution of proton to pion ratios as a function of p in the 3-jet system CM frame is shown in Fig. 25. A noticeable improvement of the ratio of junction protons to non-junction protons can be seen in the region below ~ 1 GeV. For the blind analysis, cf. Fig. 24, this ratio was about 2. The corresponding number for Fig. 25 is closer to 4, approximately the same factor as for the isolated neutralino decay, Fig. 18.

We thus conclude that even with all the surrounding activity in fully generated pp collisions, it *is* possible, with a good jet selection, to arrive at a proton sample of almost the same purity as for the case of isolated decay. It should be remembered, though, that we are not here talking about an *excess* of protons of a factor of 4, since the division into junction baryons and non-junction baryons is not an experimental observable. What is expected from ordinary $q\bar{q}$ fragmentation is represented in Fig. 18 by circles and crosshairs, and we see

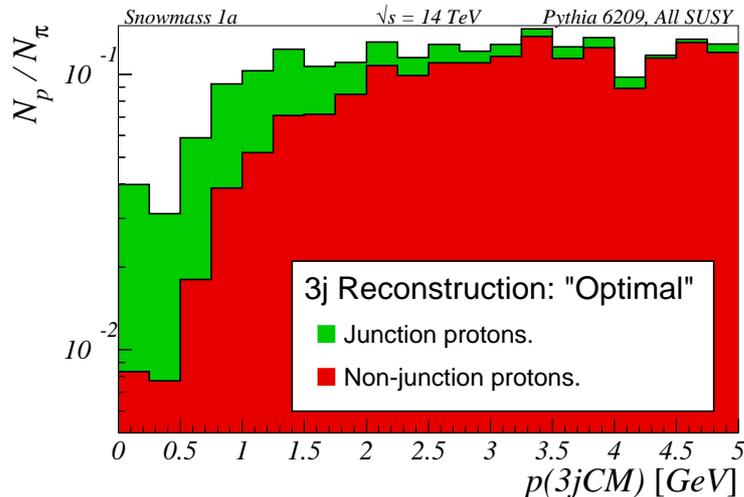


Figure 25: Reconstructed proton spectrum in the 3-jet CM frame for the “optimized” neutralino decay jet identification.

that relative to *this* number, the excess is only about a factor 1.5. We note that one does not have to rely on a Monte Carlo estimate of the proton to pion ratio. A study of reference background samples, of multijet events in configurations not significantly contaminated by SUSY events, will allow an absolute determination of the proton fraction from data. We therefore believe that even such a modest difference as a 50% enrichment should be clearly visible, given enough statistics.

Note that the rate of baryon production in high-energy collisions offers one puzzling outstanding problem in our understanding of QCD: why are fewer kaons and protons produced in $ep/\gamma p$ events than in e^+e^- ones? This problem has been with us since several years [42], and has been reconfirmed by recent HERA studies [43]. One possible explanation is the ‘quiet string scenario’ [42], namely that soft gluon emission around the perturbative–nonperturbative border is less profuse in $ep/\gamma p$ events, where only part of the event feels a hard scale, than in e^+e^- ones, where the whole event derives from a hard-scale process. The string would then be less ‘wrinkled’ in the former case than the latter, thus have a smaller effective string tension, and thereby a lesser production of heavier particles. This is very speculative, however, and having to give up jet universality would be uncomfortable, to say the least. One possible test in Deeply Inelastic Scattering would be a comparison of the current hemisphere in the Breit frame with the rest of the event. It is also unknown

whether Tevatron collider events more resemble LEP or HERA ones in this respect. In the quiet string scenario, one would expect particles associated with the perturbative jets, i.e. the particles we have studied, to resemble LEP, while the underlying event again could be quieter and therefore also contain fewer baryons. The older UA5 measurements on K_S^0 [44] do not indicate any problem for the standard string model [32]. Hopefully the issue will be settled well before a study of observed potential BNV events is initiated.

6.3 BNV two-body decays

We now follow up the discussion in subsection 5.5, considering the process $e^+e^- \rightarrow \tilde{t}_1(\rightarrow \bar{d}\bar{s})\tilde{t}_1^*(\rightarrow ds)$ for a linear collider, TESLA or the NLC/JLC, using the mSUGRA parameters of Snowmass point 5 to represent a generic “light stop” scenario. Although the 220 GeV stop mass of this benchmark point is small enough to allow pair-production at 500 GeV, the amount of junction–junction annihilation depends strongly on the stop boost, meaning that essentially no junctions would survive at CM energies close to threshold in our preferred stringlength-minimization scenario. Thus, to increase the 2-junction rate, and also to get a higher total production cross section, we are lead to consider an 800 GeV option. Note that if the BNV couplings are small, we expect string breaks between the stops to occur *before* decay. Such breaks would sever the connection in colour space between the two junctions and thus effectively eliminate the possibility of obtaining 0-junction topologies.

At the chosen CM energy, the stop boost $\gamma_{\tilde{t}} = E_{\tilde{t}}/m_{\tilde{t}} = E_{\text{CM}}/2m_{\tilde{t}} \sim 1.8$, making us expect that, in the case of large BNV couplings, about 40% of the junction–junction systems survive to give rise to junction baryons, cf. Fig. 8. The remaining 60% of the events presumably represent a more or less irreducible background to our search, since they contain no signal (junction) baryons and only differ from the signal events at the hadronization level, cf. Fig. 22 and the discussion there. Of course, once BNV is established, also the rate of junction annihilation by itself could offer interesting information on QCD.

Similarly to the LHC study, we assume that a clean sample of $\tilde{t}_1\tilde{t}_1^*$ events has been isolated, so that SM as well SUSY backgrounds can be neglected. A further simplification we make is to disregard the effects of brems- and beamstrahlung in the initial state, mostly since the latter depends sensitively on the beam parameters and we do not wish to become too machine-specific at this stage.

So as not to be completely unrealistic in regard to detector acceptance, we ignore all particles closer than 5° to the beam pipe and reconstruct only charged particles which have $p_\perp > 0.5$ GeV. Beyond these crude cuts, no effort is made to simulate detector effects. For jet reconstruction we use Durham $y_D = 10^{-2.5} \approx 0.003$. This results in about 65% 4-jet events, 20% 5-jets, and

10% 3-jets for the process considered. Naturally, the events with more than 4 jets are those which involve emissions of hard gluons. This means partly that we are facing a larger combinatorics, and partly that we are less certain in which direction the junctions went and hence in which direction to look for the junction protons. For the 3-jet events, two of the original four jets have ended up very close to each other. With a smaller y_D those jets might be resolved, but it is not clear that much could be gained by this. Thus, all but the 4-jet events are removed from the further analysis.

We begin by simply performing a proton-count within the detector acceptance, normalizing to the charged pion multiplicity, as we did for the LHC study. The result, shown in Fig. 26, is markedly better than for the LHC case, Fig. 23, chiefly due to the cleaner environment and that we now go down to $p_{\perp} = 0.5$ GeV. Even though, in Fig. 26a, we have only a 40% rate of 2-junction configurations, about 30% of the protons below 1.5 GeV are junction protons. This should be compared with Fig. 26b, where we have turned off junction–junction annihilation.

Fortunately, it is again possible to purify the sample somewhat more. First, we identify which of the 3 possible pair-by-pair combinations of the four jets is most consistent with both pairs having the invariant mass of the \tilde{t}_1 , by selecting the configuration with the smallest value of $\Delta M \equiv |m_{ij} - m_{\tilde{t}_1}| + |m_{kl} - m_{\tilde{t}_1}|$.

The mass distributions of the pairs thus selected are shown in Fig. 27 for dynamically mixed (default) 2-junction and 0-junction topologies (solid curve), for 2-junction topologies in all events (dotted curve), and for 0-junction topologies in all events (dot-dashed curve). As before, note that the latter sample has a small contamination of 2-junction events, from perturbative breakups in the $\tilde{t}_1\tilde{t}_1^*$ shower. All the distributions are shifted a few GeV towards larger masses, relative to the nominal \tilde{t}_1 mass, but the 0-junction topologies are shifted a bit more and have a wider distribution compared to the 2-junction ones. The shift of the peaks is due to a string effect: the fact that either one junction–junction string or two $q\bar{q}$ strings are spanning the range *between* the two jet systems increases the hadronic multiplicity in this region. In the former topology, all the hadrons produced in a junction–junction string add mass to whichever stop they are clustered with. (The energy of these hadrons is taken from the kinetic energy of the stops, which are slowed down correspondingly.) In the latter topology, the two separately hadronizing $q\bar{q}$ systems produce even more extra soft particles in between the \tilde{t}_1 and \tilde{t}_1^* jets. Also note that this skews the reconstructed jet directions away from the original parton directions, towards larger opening angles, which is an alternative way of expressing that the jet–jet pairs end up having invariant masses slightly above the actual resonance mass. In events with multiple gluon emissions in the $\tilde{t}_1\tilde{t}_1^*$ system — few to begin with and further reduced by our cuts — the zig-zag pattern illustrated in Fig. 7b further tends to favour large string lengths and thereby large mass shifts in the 0-junction configuration. It is important to remember that the bulk of the par-

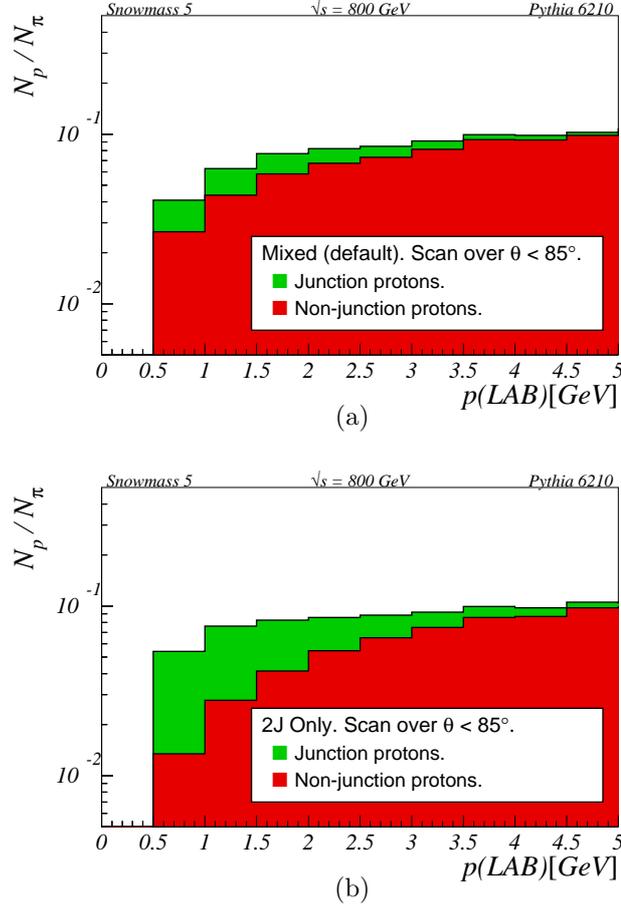


Figure 26: Reconstructed momentum spectrum for the ratio of protons to charged pions in a simple scan over the detector volume: (a) mixed 2-junction and 0-junction topologies and (b) pure 2-junction topologies.

ticle production is associated with the four (anti)quark jet directions, however, and what we discuss here are smaller perturbations on this picture.

To summarize, we do not expect that a double peak, one from the 2-junction topologies and another from the 0-junction ones, will be visible in the jet-jet mass spectrum, but we do expect a small enrichment of 2-junction topologies on the low side of the peak relative to the high side of it. Therefore we impose a slightly asymmetric cut on the jet-jet masses, $m_{\tilde{t}_1} - 10 \text{ GeV} < M_{ij} < m_{\tilde{t}_1} + 5 \text{ GeV}$. Note that both jet-jet pairs in an event are required to pass this cut, or the event is rejected, this mainly to ensure that both sides of the event are

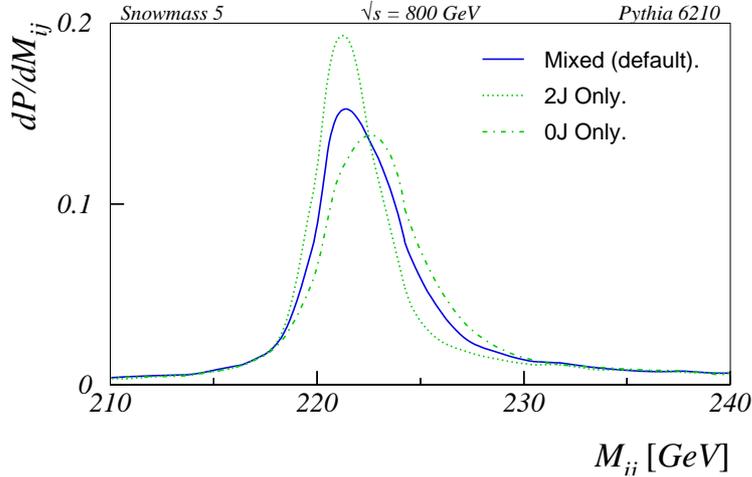


Figure 27: Reconstructed jet–jet mass spectrum. The distributions are normalized so as to integrate to the 4-jet rates.

well reconstructed.

Having thus reduced the combinatorics, it is a matter of simple kinematics to show that the two daughters from each stop decay will be separated by more than 60° at the 800 GeV CM energy. Adding parton showers and hadronization effects will not change this limit appreciably, and so we place a cut requiring jet pair opening angles larger than 60° . Larger opening angles, from asymmetric decays and/or slowed-down stops, tend to favour the 0-junction topologies, i.e. give less signal. Therefore, and since we aim to look for junction protons *between* the two daughter jets, we require a maximal jet–jet opening angle of 120° .

For the remaining jet–jet pairs, we may now safely presume that the junction baryon, if there is one, is predominantly travelling in roughly the same direction as the stop did, in between the two daughter jets. To measure how much “in between” two jets a particle is, we construct

$$\hat{\theta}_{\text{pj}_{1j_2}} = \frac{1}{\sqrt{2}} \sqrt{\left(\frac{2\theta_{\text{pj}_1}}{\theta_{j_1j_2}} - 1\right)^2 + \left(\frac{2\theta_{\text{pj}_2}}{\theta_{j_1j_2}} - 1\right)^2}, \quad (26)$$

where θ_{pj_i} is the angle between jet i and the particle in question, and $\theta_{j_1j_2}$ is the inter-jet angle. The measure is constructed so as to have the value $\hat{\theta} = 0$ when a particle is lying exactly between the two jets and $\hat{\theta} = 1$ when a particle is exactly aligned with one or the other of the jet directions. Note that $\theta_{\text{pj}_1} + \theta_{\text{pj}_2} \geq \theta_{j_1j_2}$, where the equality holds for particles lying between the two jets and in the plane defined by the jet directions.

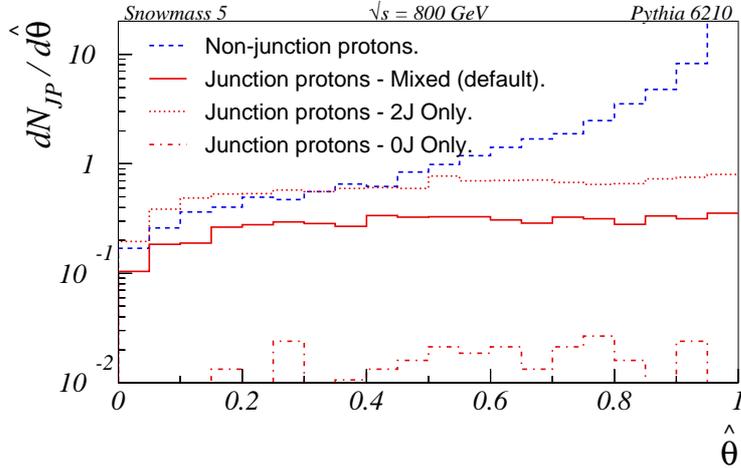


Figure 28: Distributions in $\hat{\theta}_{p_{j_1 j_2}}$ of junction protons for dynamically mixed (default) 2-junction and 0-junction topologies (solid lines), 2-junction only (dotted lines), and 0-junction only (dash-dotted lines). The dashed histogram represents the non-junction protons. These are only shown for the default case, the non-default spectra being very similar.

In Fig. 28 we show the distributions of junction protons and ordinary protons in this variable for events passing the previous cuts. Note that the junction proton distributions, even for the pure 2-junction topologies, do not integrate to 2 over the $\hat{\theta}$ range shown on the plot. This owes to the simple facts that many junction protons are irretrievably lost by the $p_{\perp} \geq 0.5$ GeV cut we imposed above and that a non-vanishing fraction of them have $\hat{\theta}_{p_{j_1 j_2}} > 1$. The logarithmic scale on the plot is called for by the very strong peaking of the non-junction proton spectrum towards the jet axes. We observe that the shape of the junction proton spectrum is rather flat, independently of whether we allow junction–junction annihilation or not, although the normalization of course changes. When junction–junction annihilation is forced (dot-dashed curve), only those junctions which are separated by $g \rightarrow q\bar{q}$ splittings in the cascades of the $\tilde{t}_1\tilde{t}_1^*$ system survive after hadronization, hence extremely few junction protons are produced. We impose a cut requiring $\hat{\theta}_{p_{j_1 j_2}} < 0.5$, thereby cleaning out most of the jet-associated protons.

Furthermore, we may use the fact that the junction proton momentum spectrum is most strongly peaked in the junction rest frame (JRF), cf. Fig. 17. Analogously to what we did for the LHC study, we therefore construct an approximate JRF for each jet pair, in each case using the linear combination of the momenta of the opposing pair as the “third leg”. Including the above-

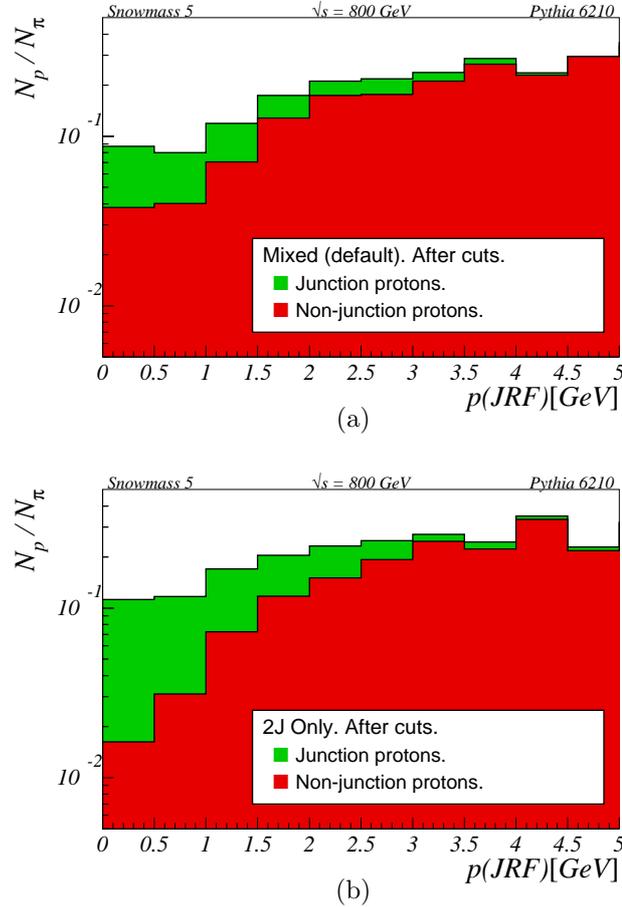


Figure 29: Reconstructed momentum spectrum in the approximated junction rest frame for the ratio of protons to charged pions: (a) dynamically mixed (default) 2-junction and 0-junction topologies and (b) 2-junction topologies only.

mentioned cuts, this results in the multiplicity distribution Fig. 29a, now with about 60% of the protons below 1 GeV being junction protons. A rough estimate is that this corresponds to about 50 junction protons per 100 fb^{-1} of integrated luminosity. With junction–junction annihilation turned off, the distribution changes to that of Fig. 29b, with about 80% of the protons below 1 GeV being junction protons, corresponding to more than 100 junction protons per 100 fb^{-1} .

Thus, for the process and SUSY scenario considered here, small BNV couplings substantially increase the possibility of directly detecting the extra (anti)baryons.

Going further and establishing the absolute rate of junction protons per $\tilde{t}_1\tilde{t}_1^*$ event, with the BNV branching fractions known, would obviously amount to a direct measurement of the junction–junction annihilation rate. Such a measurement would be interesting from a QCD perspective, since it enables us to confront the value predicted by our model with an experimental result. We therefore here also briefly discuss an alternative way of measuring this quantity.

In the 0-junction topology, with at least two strings spanning the rapidity region between the stops, we here expect a higher average multiplicity than for the 2-junction case, in which only one string spans this region. Naïvely then, the charged multiplicity between the two jet pairs should be at least twice as large for 0-junction topologies as compared to 2-junction ones. However, this drastically oversimplifies the true situation. Asymmetric decays of the stops and the addition of parton showers will in general smear the difference to a much lower factor, to the point where it becomes questionable whether this method is at all feasible. To investigate, we select events passing the above-mentioned cuts and plot the charged multiplicity as a function of the absolute value of the rapidity with respect to the stop–antistop momentum axis. The resulting rapidity distributions, shown in Fig. 30a, only exhibit a negligible increase in the ratio $(dN_{0j}/d|y|)/(dN_{2j}/d|y|)$ in the region $|y| \sim 0$.

One reason for this is that the central region is not free from jet activity. Specifically, if the decay products of one of the stops are aligned with the original direction of motion of the stop, the backward jet will end up in the “wrong hemisphere”. The most extreme of these cases are already removed by the requirement of jet–jet opening angles smaller than 120° , as was imposed above, but this cut still leaves room for jet activity in the mid-rapidity region. To clean out the worst of this contamination, in Fig. 30b, we have tightened the cut on maximal jet–jet opening angles to $\theta_{jj} \leq 90^\circ$. One notices that the average multiplicity in the central region drops, and that the ratio between the multiplicities in the two topologies increases in this region. Ultimately, to get as little contamination from jet tails as possible in the central region, one should go as close as possible to the kinematic limit, $\theta_{jj} = 60^\circ$, where the rapidity of all jets, with respect to the stop–antistop momentum axis, is $|y| \sim 1.3$. The necessary compromise thus lies between minimizing the jet contamination in the central region and obtaining sufficient statistics. For the present, we merely note that the central multiplicity *is* sensitive to the choice of string topology but that it will presumably be hard to use it for a precise determination of the junction–junction annihilation rate.

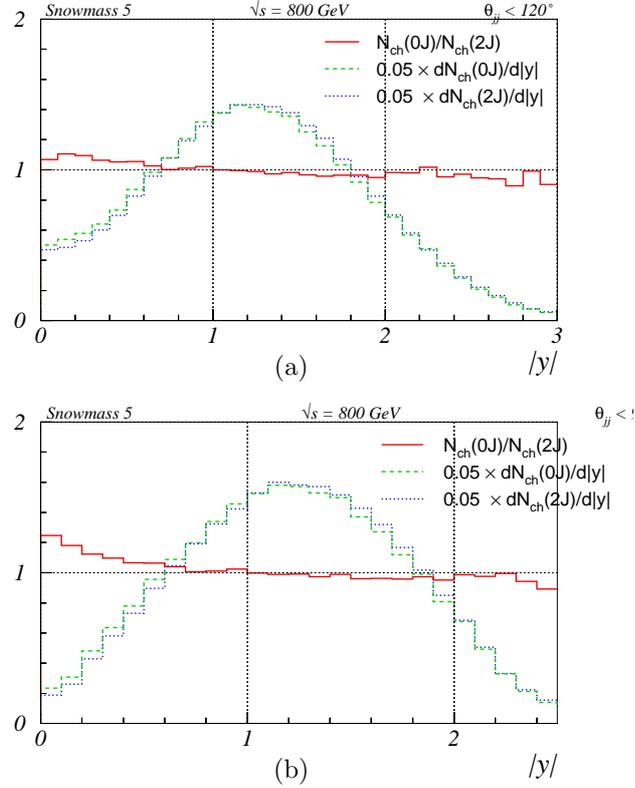


Figure 30: Charged multiplicities of 0-junction topologies (dashed lines) and 2-junction topologies (dotted lines) as functions of rapidity along the stop-antistop thrust axis for (a) $\theta_{jj} \leq 120^\circ$ and (b) $\theta_{jj} \leq 90^\circ$ (all other cuts are as for the junction proton search). The ratios of the distributions are shown with solid lines.

7 Summary

The physics scenario studied here may not be the most likely one. First we have to accept supersymmetry, then that R -parity is violated, and finally that it is baryon rather than lepton number that is affected. Nevertheless, it is one possibility among many that should be considered in studies of the high-energy frontier, and one that has maybe received less attention than others of equal likelihood to be correct.

In particular, it is important to know whether such a scenario introduces new special demands on detectors. Obviously the whole normal artillery of jet

reconstruction and mass peaks need to be brought to bear on the issue, in order to find a signal for new physics. This may already be enough to favour some specific scenarios, such as BNV SUSY. However, given that the unique aspect of BNV is the creation of an additional baryon or antibaryon, a discovery of BNV SUSY would not be complete without a clear signal that indeed baryon number is violated.

Since baryons are produced also in normal hadronization processes, it will not be feasible to associate one specific baryon uniquely with the BNV process, but only to establish an excess of baryon production. This excess will not be large enough to help us, unless we know where to look, and thereby can reduce the background noise level.

The string junction model we develop in this article offers us the tools to do precisely that. The string concept itself is the most successful approach we have to the description of hadronization at high energies. The junction concept is also well-established in the literature, to describe baryon colour topologies and confinement aspects. It has never before (to the best of our knowledge) been developed and implemented into a realistic hadronization framework, however.

Armed with this detailed model, we can predict where the baryon excess should be found. In general, the junction baryon will have low momentum in the rest frame of the process that generates the BNV. For the specific case of neutralino or chargino BNV decay at LHC, we show how this translates into explicit predictions. Typically we need to impose some cut like $p_{\perp} > 1$ GeV to get away from the bulk of the underlying event activity, but beyond that the main demand is for a good p/\bar{p} identification at rather low momenta, i.e. a few GeV. Failing this, at least Λ tagging should be well developed.

We believe this to be the first example where non-QCD physics interests put precise demands on baryon identification. Thereby it offers an interesting test case for further detector performance studies. The specific distributions we show for LHC are only intended to “whet the appetite”, and are in no way blueprints for a realistic analysis. Background studies will be essential, not only in the search of a signal, but also to help calibrate the normal rate of baryon production. Especially disconcerting is here the experimental indications of a difference in the baryon fraction between different kinds of events, e.g. ep vs. e^+e^- .

Studies could also be performed for other machines. Unfortunately, it is doubtful whether the Tevatron could produce enough BNV events to allow an analysis along the lines we have in mind, given that no signs of SUSY have been seen so far. In contrast, BNV neutralino or chargino decays at a future linear e^+e^- collider should be much cleaner and thereby simpler to analyse — one reason we have not considered it here. The difficult questions at such a collider could rather come in the analysis of BNV $\tilde{t}\tilde{t}^*$ events, where the threshold region production would allow two different event topologies, one with an extra baryon–antibaryon pair and one without. We find that the relative importance

of these topologies may not be easy to pin down experimentally.

It can be questioned how unique the proposed junction hadronization framework is, given that it is not based on any experimental tests to existing data. On the technical level, certainly a number of approximations and simplifications have been introduced, but we have demonstrated that the resulting uncertainties appear under control. More serious is the issue whether the whole ansatz as such is what happens in Nature, and here no guarantees can be given. What can be said is that there is currently no alternative approach that looks anywhere near as credible.

One specific example is here offered by HERWIG. This generator maps a three-quark cluster onto a quark–diquark one, which then is fragmented along a single axis. Since the baryon cluster mass typically is very large — in part a consequence of a low shower activity in these events — the fragmentation can deform the event appreciably and, more significantly, kick the baryon out to surprisingly large momenta.

That is not to say there is no room for improvement in our PYTHIA implementation. We did not (yet) include the matrix-element information to give non-isotropic three-body BNV decays, the BNV production processes are not implemented at all, and our shower description is not as sophisticated as e.g. for $Z^0 \rightarrow q\bar{q}$ decays. Such imperfections could influence estimates of the experimental acceptance rate of BNV decays. Given all the other uncertainties at this stage of the game, one should not exaggerate their importance, however. Once SUSY is observed, with indications of BNV, further work would be reasonable.

The development of a model for the hadronization of junction string topologies is not only relevant for BNV SUSY. In principle we would assume the same topology for the incoming baryons in high-energy pp or $p\bar{p}$ collisions. So long as only one valence quark is kicked out of a proton, the two remaining quarks drag the junction along and thereby form an effective diquark. But in multiple interaction approaches there is every possibility for two quarks to be kicked out, in different directions, and with colours rearranged in the process. Thereby the baryon number can start to drift in the event. This scenario was advertised in [32], but has not been studied further till now. The recent increased interest in semihard physics at hadron colliders [45], which has confirmed the basic validity of the PYTHIA framework but also pointed at problems, here provides a stimulus for further developments. Hopefully, refined models could allow us to understand the excess of baryon number observed in the central region of events in a number of experiments ([46] is a far from complete list). Here junction scenarios have already been proposed as a possible mechanism [47], but without any solid modelling efforts.

A reason for us to begin with the SUSY BNV studies, rather than with pure QCD ones, is that the latter inevitably will come to involve further uncertainties, such as the details of the multiple interactions scenario. Nevertheless, the

hope is that continued QCD-related studies may add support for the validity of the junction scenario proposed in this article. Thereby the physics of junction fragmentation offers a prime example of a topic where the exploration of very conventional and very unconventional physics goes hand in hand.

Acknowledgements

The authors would like to thank Gösta Gustafson for helpful discussions on the string length of junction topologies, and Frank Paige and Albert de Roeck for useful information on the ATLAS and CMS proton identification capabilities. We are also grateful to the NorduGrid project, for inviting us to test their facilities. Many of our results were obtained with grid computing.

References

- [1] I. Hinchliffe and T. A. Kaeding, *Phys. Rev. D* **54** (1996) 914.
- [2] H. Baer, F.E. Paige, S.D. Protopopescu and X. Tata, hep-ph/0001086.
- [3] G. Corcella, I.G. Knowles, G. Marchesini, S. Moretti, K. Odagiri, P. Richardson, M.H. Seymour and B.R. Webber, *JHEP* **0101** (2001) 010
- [4] S. Katsanevas and P. Morawitz, *Comput. Phys. Commun.* **112** (1998) 227.
- [5] T. Sjöstrand, P. Edén, C. Friberg, L. Lönnblad, G. Miu, S. Mrenna and E. Norrbin, *Computer Phys. Commun.* **135** (2001) 238;
T. Sjöstrand, L. Lönnblad and S. Mrenna, LU TP 01-21 [hep-ph/0108264].
- [6] P.Z. Skands, *Cand. Scient. thesis*, Copenhagen University (2001), hep-ph/0108207.
- [7] P.Z. Skands, *Eur. Phys. J.* **C23** (2002) 173.
- [8] H.K. Dreiner, P. Richardson and M.H. Seymour, *JHEP* **0004** (2000) 008.
- [9] S. Moretti, K. Odagiri, P. Richardson, M.H. Seymour and B.R. Webber, *JHEP* **0204** (2002) 028.
- [10] B.C. Allanach, A.J. Barr, L. Drage, C.G. Lester, D. Morgan, M.A. Parker, B.R. Webber and P. Richardson, *JHEP* **0103** (2001) 0248;
B.C. Allanach, A.J. Barr, M.A. Parker, P. Richardson and B.R. Webber, *JHEP* **0109** (2001) 021.
- [11] P. Z. Skands, to appear in the proceedings of SUSY02, Hamburg, Germany, 17-23 Jun 2002, LU TP 02-35 [hep-ph/0209199].
- [12] X. Tata, Lectures presented at the IX JAS Summer School, hep-ph/9706307.
- [13] S. Dimopoulos et al., *Phys. Rev. D* **41** (1990) 2099;
R. J. Oakes et al., *Phys. Rev. D* **57** (1998) 534;
B. Allanach et al., contributed to ‘Physics at Run II: Workshop on Supersymmetry/Higgs’ (1998), hep-ph/9906224;
E. L. Berger, B. W. Harris and Z. Sullivan, *Phys. Rev. D* **63** (2001) 115001.

- [14] L. Lönnblad, *Computer Phys. Commun.* **71** (1992) 15.
- [15] I.G. Knowles et al., in ‘Physics at LEP2’, eds. G. Altarell, T. Sjöstrand and F. Zwirner, CERN 96–01, Vol. 2, p. 103;
P. Bambade et al., in ‘Reports of the Working Groups on Precision Calculations for LEP2 Physics’, eds. S. Jadach, G. Passarino and R. Pittau, CERN 2000–009, p. 137.
- [16] M.H. Seymour, *Computer Phys. Commun.* **90** (1995) 95.
- [17] G. Corcella and M.H. Seymour, *Phys. Lett.* **B442** (1998) 417.
- [18] M. Bengtsson and T. Sjöstrand, *Phys. Lett.* **B185** (1987) 435; *Nucl. Phys.* **B289** (1987) 810
- [19] E. Norrbin and T. Sjöstrand, *Nucl. Phys.* **B603** (2001) 297.
- [20] V.A. Khoze, L.H. Orr and W.J. Stirling, *Nucl. Phys.* **B378** (1992) 413.
- [21] B. Andersson, G. Gustafson, G. Ingelman and T. Sjöstrand, *Phys. Rep.* **97** (1983) 31;
B. Andersson, ‘The Lund Model’ (Cambridge University Press, 1998).
- [22] R.K. Ellis, W.J. Stirling and B.R. Webber, ‘QCD and Collider Physics’ (Cambridge University Press, 1996), and references therein.
- [23] B. Andersson, G. Gustafson and T. Sjöstrand, *Nucl. Phys.* **B197** (1982) 45.
- [24] B. Andersson, G. Gustafson and T. Sjöstrand, *Physica Scripta* **32** (1985) 574;
P. Edén and G. Gustafson, *Z. Phys.* **C75** (1997) 41.
- [25] I.G. Knowles and G.D. Lafferty, *J. Phys.* **G23** (1997) 731, and references therein.
- [26] T. Sjöstrand, *Nucl. Phys.* **B248** (1984) 469.
- [27] X. Artru, *Nucl. Phys.* **B85** (1975) 442.
- [28] M. Imachi, S. Otsuki and F. Toyoda, *Progr. Theor. Phys.* **54** (1975) 280;
Y. Igarashi et al., *Progr. Theor. Phys. Suppl.* **63** (1978) 49;
G.C. Rossi and G. Veneziano, *Nucl. Phys.* **B123** (1977) 507.
- [29] L. Montanet, G.C. Rossi and G. Veneziano, *Phys. Rep.* **63** (1980) 149.
- [30] G. ’t Hooft, *Physica Scripta* **25** (1982) 133.
- [31] B. Andersson, G. Gustafson, I. Holgerdsson and O. Månsson, *Nucl. Phys.* **B178** (1981) 242.
- [32] T. Sjöstrand and M. van Zijl, *Phys. Rev.* **D36** (1987) 2019.
- [33] I. Montvay, *Phys. Lett.* **84B** (1979) 331.
- [34] V.A. Khoze and T. Sjöstrand, *Phys. Lett.* **B328** (1994) 466.
- [35] B. Andersson, G. Gustafson and B. Söderberg, *Z. Phys.* **C20** (1983) 317;
B. Andersson, S. Mohanty and F. Söderberg, *Nucl. Phys.* **B646** (2002) 102.

-
- [36] B. C. Allanach et al., in ‘Proc. of the APS/DPF/DPB Summer Study on the Future of Particle Physics’ (Snowmass 2001), eds. R. Davidson and C. Quigg [hep-ph/0202233].
- [37] S. Catani, Yu.L. Dokshitzer, M. Olsson, G. Turnock and B.R. Webber, *Phys. Lett.* **B269** (1991) 432.
- [38] S. Moretti, L. Lönnblad and T. Sjöstrand, *JHEP* **9808** (1998) 001.
- [39] R.K. Ellis, D.A. Ross, and A.E. Terrano, *Nucl. Phys.* **B148** (1981) 421.
- [40] ATLAS Collaboration, Technical Report, CERN/LHCC 99-14.
- [41] Particle Data Group, K. Hagiwara et al., *Phys. Rev.* **D66** (2002) 010001.
- [42] T. Sjöstrand, *J. Phys.* **G22** (1996) 709, and references therein.
- [43] H1 Collaboration, abstract 1002, submitted to ICHEP02, Amsterdam, The Netherlands, 24–31 July 2002.
- [44] UA5 Collaboration, G.J. Alner et al., *Nucl. Phys.* **B258** (1985) 505.
- [45] CDF Collaboration, D. Acosta et al., *Phys. Rev.* **D65** (2002) 072005;
CDF Collaboration, T. Affolder et al., *Phys. Rev.* **D65** (2002) 092002;
D0 Collaboration, V.M. Abazov et al., hep-ex/0207046;
R.D. Field and A. Kharchilava, presentations at the ‘Matrix Element and Monte Carlo Tuning Workshop’, Fermilab, October 4, 2002,
<http://www-cpd.fnal.gov/MCTuning/>.
- [46] G. Belletini et al., *Nuovo Cimento* **A42** (1977) 85;
B. Alper et al., *Nucl. Phys.* **B100** (1975) 237;
A. Breakstone et al., *Z. Physik* **C28** (1985) 335;
H1 Collaboration, paper 556, submitted to ICHEP98, Vancouver, July 1998, and in preparation;
STAR Collaboration, C. Adler et al., *Phys. Rev. Lett.* **86** (2001) 4778;
PHENIX Collaboration, K. Adcox et al., *Phys. Rev. Lett.* **89** (2002) 092302;
PHOBOS Collaboration, B.B. Back et al., nucl-ex/0206012;
BRAHMS Collaboration, I.G. Bearden et al., nucl-ex/0207006.
- [47] D. Kharzeev, *Phys. Lett.* **B378** (1996) 238;
B. Kopeliovich and B. Povh, *Z. Physik* **C75** (1997) 693;
S.E. Vance, M. Gyulassy and X.-N. Wang, *Phys. Lett.* **B443** (1998) 45;
G.T. Garvey, B.Z. Kopeliovich, B. Povh, *Comments Mod. Phys.* **A2** (2001) 47.

**SUSY Les Houches Accord:
Interfacing SUSY Spectrum
Calculators, Decay Packages, and
Event Generators**

Paper III

SUSY Les Houches Accord: Interfacing SUSY Spectrum Calculators, Decay Packages, and Event Generators

P. Skands¹, B.C. Allanach², H. Baer³, C. Balázs^{3,4}, G. Bélanger²,
F. Boudjema², A. Djouadi^{5,6}, R. Godbole⁷, J. Guasch⁸, S. Heinemeyer^{6,9},
W. Kilian¹⁰, J-L. Kneur⁵, S. Kraml⁶, F. Moortgat¹¹, S. Moretti¹²,
M. Mühlleitner⁸, W. Porod¹³, A. Pukhov¹⁴, P. Richardson^{6,15},
S. Schumann¹⁶, P. Slavich¹⁷, M. Spira⁸, G. Weiglein¹⁵

¹ *Theoretical Physics, Lund University, Sölvegatan 14A, 22362 Lund, Sweden.*

² *LAPTH, 9 Chemin de Bellevue, BP 110 74941 Annecy-le-Vieux, Cedex, France.*

³ *Dept. of Physics, Florida State University, Tallahassee, FL 32306, U.S.A.*

⁴ *HEP Division,*

Argonne National Laboratory, 9700 Cass Ave., Argonne, IL 60439, U.S.A.

⁵ *LPMT, Université Montpellier II, 34095 Montpellier, Cedex 5, France.*

⁶ *Theory Division, CERN, CH-1211 Geneva 23, Switzerland.*

⁷ *CTS, Indian Institute of Science, Bangalore, 560012, India.*

⁸ *Paul Scherrer Institut, CH-5232 Villigen PSI, Switzerland.*

⁹ *LMU München, Theresienstr. 37, D-80333 München, Germany.*

¹⁰ *Theory Group, DESY, D-22603 Hamburg, Germany.*

¹¹ *University of Antwerpen, B-2610 Antwerpen, Belgium.*

¹² *School of Physics and Astronomy,*

University of Southampton, Highfield, Southampton SO17 1BJ, UK.

¹³ *Institute for Theoretical Physics,*

University of Zürich, Winterthurerstr. 190, CH-8057 Zürich, Switzerland.

¹⁴ *Moscow State University, Russia.*

¹⁵ *IPPP, University of Durham, Durham DH1 3LE, UK.*

¹⁶ *Institut für Theoretische Physik, TU Dresden, D-01062 Dresden, Germany.*

¹⁷ *Max Planck Institut für Physik, Föhringer Ring 6, D-80805 München, Germany*

Abstract An accord specifying a unique set of conventions for supersymmetric extensions of the Standard Model together with generic file structures for 1) supersymmetric model specifications and input parameters, 2) electroweak scale supersymmetric mass and coupling spectra, and 3) decay tables is presented, to provide a universal interface between spectrum calculation programs, decay packages, and high energy physics event generators.

¹Peter.Skands@thep.lu.se. See www.thep.lu.se/~zeiler/slha/ for updates and examples.

1 Introduction

An increasing number of advanced programs for the calculation of the supersymmetric (SUSY) mass and coupling spectrum are appearing [1–5] in step with the more and more refined approaches which are taken in the literature. Furthermore, these programs are often interfaced to specialized decay packages [4–8], relic density calculations [9, 10], and (parton-level) event generators [11–18], in themselves fields with a proliferation of philosophies and, consequentially, programs.

At present, a small number of specialized interfaces exist between various codes. Such tailor-made interfaces are not easily generalized and are time-consuming to construct and test for each specific implementation. A universal interface would clearly be an advantage here.

However, since the codes involved are not all written in the same programming language, the question naturally arises how to make such an interface work across languages. At present, an inter-language linking solution does not seem to be feasible without introducing at least some dependence on platform (e.g. UNIX variant) and/or compiler. For details on these aspects, see e.g. [19].

At this point, we deem such an interface too fragile to be set loose among the particle physics community. Instead, we advocate a less elegant but more robust solution, exchanging information between FORTRAN and C(++) codes via three ASCII files, one for model input, one for model input plus spectrum output, and one for model input plus spectrum output plus decay information. The detailed structure of these files is described in the sections below. Briefly stated, the purpose of this Accord is thus the following:

1. To present a set of generic definitions for an input/output file structure which provides a universal framework for interfacing SUSY spectrum calculation programs.
2. To present a generic file structure for the transfer of decay information between decay calculation packages and event generators.

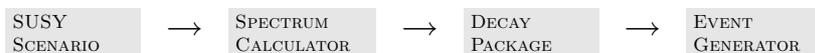


Figure 1: Stages of the interface Accord. By SUSY scenario we not only intend specific SUSY-breaking mechanisms, such as supergravity (SUGRA), gauge mediated SUSY breaking (GMSB), etc., but also more general setups within the Minimal Supersymmetric Standard Model (MSSM).

Note that different codes may have different implementations of how SUSY Les Houches Accord (SLHA) input/output is *technically* achieved. The details of

how to ‘switch on’ SLHA input/output with a particular program should be described in the manual of that program and are not covered here.

1.1 Using the Accord

To interface two or more calculations, the general procedure would be that the user prepares model input parameters together with a set of Standard Model parameters (to be used as low-scale boundary conditions for the spectrum calculation) in an ASCII file, complying with the standard defined in section 3.1 below. At present, only models with the particle spectrum of the MSSM, and with CP and R-parity conserved are included in this standard.

The user then runs a spectrum calculation program with these inputs to obtain the SUSY mass and coupling spectrum at the electroweak (EW) scale. The resulting spectrum is stored, together with a copy of the model input parameters (so that subsequent calculations may be performed consistently), in a new file. Standards for the spectrum file are defined in section 3.2.

The user may now run some particular decay package to generate a list of decay modes and widths for selected particles, which the decay package saves to a third file, complying with the definitions in section 3.3. Again, a copy of the model input parameters as well as the complete spectrum information is included together with the decay information in this file.

Lastly, the user may instruct a (parton-level) event generator to read in all this information and start generating events. Of course, any of these intermediate steps may be skipped whenever the user does not wish to switch between programs across them, e.g. no decay information is required to be present in the file read by an event generator if the user wishes the event generator to calculate all decay widths itself.

If a general purpose event generator is used, the events will include parton showering and hadronization, whereas if a parton-level generator is used, the events may finally be passed to parton showering and hadronization programs using the already defined Les Houches Accord #1 [20].

2 Conventions

One aspect of supersymmetric calculations that has often given rise to confusion and consequent inconsistencies in the past is the multitude of ways in which the parameters can be, and are being, defined. Hoping to minimize both the extent and impact of such confusion, we have chosen to adopt one specific set of self-consistent conventions for the parameters appearing in this Accord. These conventions are described in the following subsections. As yet, we only consider R-parity and CP conserving scenarios, with the particle spectrum of the MSSM.

Compared to the widely used Gunion and Haber conventions and notation [21], our prescriptions show a few differences. These will be remarked upon in all places where they occur below, with parameters in the notations and conventions of [21] denoted with an explicit superscript: ^{GH}.

2.1 Standard Model Parameters

In general, the SUSY spectrum calculations impose low-scale boundary conditions on the renormalization group equation (RGE) flows to ensure that the theory gives correct predictions for low-energy observables. Thus, experimental measurements of masses and coupling constants at the electroweak scale enter as inputs to the spectrum calculators.

In this Accord, we choose a specific set of low-scale input parameters (block SMINPUTS below), letting the electroweak sector be fixed by

1. $\alpha_{\text{em}}(m_Z)^{\overline{\text{MS}}}$: the electromagnetic coupling at the Z pole in the $\overline{\text{MS}}$ scheme with 5 active flavours (see e.g. [22]¹). This coupling is connected to the classical fine structure constant, $\alpha = 1/137.0359895(61)$ [22] by a relation of the form:

$$\alpha_{\text{em}}(m_Z)^{\overline{\text{MS}}} = \frac{\alpha}{1 - \Delta\alpha(m_Z)^{\overline{\text{MS}}}}, \quad (1)$$

where $\Delta\alpha(m_Z)^{\overline{\text{MS}}}$ contains the quantum corrections involved in going from the classical limit to the $\overline{\text{MS}}$ value at the scale m_Z .

2. G_F : the Fermi constant determined from muon decay.
3. m_Z : the Z boson pole mass.

All other electroweak parameters, such as m_W and $\sin^2\theta_W$, should be derived from these inputs if needed.

The strong interaction strength is fixed by $\alpha_s(m_Z)^{\overline{\text{MS}}}$ (the five-flavour strong coupling at the scale m_Z in the $\overline{\text{MS}}$ scheme), and the third generation Yukawa couplings are obtained from the top and tau pole masses, and from $m_b(m_b)^{\overline{\text{MS}}}$, see [22]. The reason we take the running b mass in the $\overline{\text{MS}}$ scheme rather than a pole mass definition is that the latter suffers from infra-red sensitivity problems, hence the former is the quantity which can be most accurately related to experimental measurements. If required, relations between running and pole quark masses may be found in [23, 24].

It is also important to note that all presently available experimental determinations of e.g. α_s and the running b mass are based on assuming the Standard Model as the underlying theory, for natural reasons. When extending the field content of the SM to that of the MSSM, the *same* measured results would be obtained for *different* values of these quantities, due to the

¹Note that [22] uses the notation $\hat{\alpha}(M_Z)$ for this parameter.

different underlying field content present in the MSSM. However, since these values are not known before the spectrum has been determined, all parameters contained in block `SMINPUTS` should be the ‘ordinary’ ones obtained from SM fits, i.e. with no SUSY corrections included. The spectrum calculators themselves are then assumed to convert these parameters into ones appropriate to an MSSM framework.

Finally, while we assume $\overline{\text{MS}}$ running quantities with the SM as the underlying theory as input, all running parameters in the *output* of the spectrum calculations are defined in the modified dimensional reduction ($\overline{\text{DR}}$) scheme [25–27], with different spectrum calculators possibly using different prescriptions for the underlying effective field content. More on this in section 2.5.

2.2 Supersymmetric Parameters

The chiral superfields of the MSSM have the following $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$ quantum numbers

$$\begin{aligned} L : & (1, 2, -\tfrac{1}{2}), & \bar{E} : & (1, 1, 1), & Q : & (3, 2, \tfrac{1}{6}), & \bar{U} : & (\bar{3}, 1, -\tfrac{2}{3}), \\ \bar{D} : & (\bar{3}, 1, \tfrac{1}{3}), & H_1 : & (1, 2, -\tfrac{1}{2}), & H_2 : & (1, 2, \tfrac{1}{2}). \end{aligned} \quad (2)$$

Then, the superpotential (omitting RPV terms) is written as

$$W = \epsilon_{ab} [(Y_E)_{ij} H_1^a L_i^b \bar{E}_j + (Y_D)_{ij} H_1^a Q_i^b \bar{D}_j + (Y_U)_{ij} H_2^b Q_i^a \bar{U}_j - \mu H_1^a H_2^b]. \quad (3)$$

Throughout this section, we denote $SU(2)_L$ fundamental representation indices by $a, b = 1, 2$ and generation indices by $i, j = 1, 2, 3$. Colour indices are everywhere suppressed, since only trivial contractions are involved. ϵ_{ab} is the totally antisymmetric tensor, with $\epsilon_{12} = \epsilon^{12} = 1$. Lastly, we will use t, b, τ to denote the $i = j = 3$ entries of mass or coupling matrices (top, bottom and tau).

The Higgs vacuum expectation values (VEVs) are $\langle H_i^0 \rangle = v_i/\sqrt{2}$, and $\tan \beta = v_2/v_1$. We also use the notation $v = \sqrt{v_1^2 + v_2^2}$. Different choices of renormalization scheme and scale are possible for defining $\tan \beta$. For the input to the spectrum calculators, we adopt by default the commonly encountered definition

$$\tan \beta(m_Z)^{\overline{\text{DR}}}, \quad (4)$$

i.e. the $\tan \beta$ appearing in block `MINPAR` below is defined as a $\overline{\text{DR}}$ running parameter given at the scale m_Z . The optional extended input block `EXTPAR` allows the possibility of using an input definition at a different scale, $\tan \beta(M_{\text{input}} \neq m_Z)^{\overline{\text{DR}}}$. Lastly, the spectrum calculator may be instructed to write out one or several values of $\tan \beta(Q)^{\overline{\text{DR}}}$ at various scales Q_i , see section 2.5 and block `HMIX` below.

The MSSM $\overline{\text{DR}}$ gauge couplings (block **GAUGE** below) are: g' (hypercharge gauge coupling in Standard Model normalization), g ($SU(2)_L$ gauge coupling) and g_3 (QCD gauge coupling).

Our Yukawa matrices, Y_E , Y_D , and Y_U , correspond exactly to $(f)^{\text{GH}}$, $(f_1)^{\text{GH}}$, and $(f_2)^{\text{GH}}$, respectively, in the notation of [21]. For hypercharge, [21] uses $(y)^{\text{GH}} \equiv 2Y$, and for the $SU(2)_L$ singlet leptonic superfield the notation $(\hat{R})^{\text{GH}} \equiv \bar{E}$. Finally, for the Higgs vacuum expectation values, we choose the convention in which $\langle H_i^0 \rangle = v_i/\sqrt{2} \equiv (v_i)^{\text{GH}}$, so that $v^2 = (v_1^2 + v_2^2) = (246 \text{ GeV})^2$, corresponding to $m_Z^2 = \frac{1}{4}(g'^2 + g^2)(v_1^2 + v_2^2)$, whereas [21] has $(v_1^2 + v_2^2)^{\text{GH}} = (174 \text{ GeV})^2$, with $m_Z^2 = \frac{1}{2}(g'^2 + g^2)(v_1^2 + v_2^2)^{\text{GH}}$. Otherwise, conventions for the superpotential are identical between this article and [21]².

2.3 SUSY Breaking Parameters

We now tabulate the notation of the soft SUSY breaking parameters. The trilinear scalar interaction potential is

$$V_3 = \epsilon_{ab} \sum_{ij} \left[(T_E)_{ij} H_1^a \tilde{L}_{iL}^b \tilde{e}_{jR}^* + (T_D)_{ij} H_1^a \tilde{Q}_{iL}^b \tilde{d}_{jR}^* + (T_U)_{ij} H_2^b \tilde{Q}_{iL}^a \tilde{u}_{jR}^* \right] + \text{h.c.} , \quad (5)$$

where fields with a tilde are the scalar components of the superfield with the identical capital letter. In the literature the T matrices are often decomposed as

$$\frac{T_{ij}}{Y_{ij}} = A_{ij} \quad ; (\text{no sum over } i, j) , \quad (6)$$

where Y are the Yukawa matrices and A the soft supersymmetry breaking trilinear couplings. See also blocks **YE**, **YD**, **YU**, **AE**, **AD**, and **AU** below.

The scalar bilinear SUSY breaking terms are contained in the potential

$$V_2 = m_{H_1}^2 H_{1a}^* H_1^a + m_{H_2}^2 H_{2a}^* H_2^a + \tilde{Q}_{iLa}^* (m_{\tilde{Q}}^2)_{ij} \tilde{Q}_{jL}^a + \tilde{L}_{iLa}^* (m_{\tilde{L}}^2)_{ij} \tilde{L}_{jL}^a + \tilde{u}_{iR} (m_{\tilde{u}}^2)_{ij} \tilde{u}_{jR}^* + \tilde{d}_{iR} (m_{\tilde{d}}^2)_{ij} \tilde{d}_{jR}^* + \tilde{e}_{iR} (m_{\tilde{e}}^2)_{ij} \tilde{e}_{jR}^* - (m_3^2 \epsilon_{ab} H_1^a H_2^b + \text{h.c.}) . \quad (7)$$

Rather than using m_3^2 itself, below we use the more convenient parameter m_A^2 , defined by:

$$m_A^2 = \frac{m_3^2}{\sin \beta \cos \beta} , \quad (8)$$

which is identical to the pseudoscalar Higgs mass at tree level in our conventions.

²The sign of μ in the original Gunion and Haber article actually disagrees with ours, but that sign was subsequently changed in the erratum to that article, which we here include when giving reference to [21].

Writing the bino as \tilde{b} , the unbroken $SU(2)_L$ gauginos as $\tilde{w}^{A=1,2,3}$, and the gluinos as $\tilde{g}^{X=1\dots 8}$, the gaugino mass terms (appearing in blocks EXTPAR and GAUGE below) are contained in the Lagrangian

$$\mathcal{L}_G = \frac{1}{2} \left(M_1 \tilde{b}\tilde{b} + M_2 \tilde{w}^A \tilde{w}^A + M_3 \tilde{g}^X \tilde{g}^X \right) + \text{h.c.} . \quad (9)$$

For the soft trilinear breaking terms above, we use the same sign convention as [21], but with a different normalization and unit $(m_6 A_i)^{\text{GH}} \equiv A_i$. For the bilinear breaking terms, we differ only in notation; $(m_i)^{\text{GH}} \equiv m_{H_i}$, $(m_{12})^{\text{GH}} \equiv m_3$, $(M')^{\text{GH}} \equiv M_1$, $(M)^{\text{GH}} \equiv M_2$, $(\lambda')^{\text{GH}} \equiv \tilde{b}$, and $(\lambda^A)^{\text{GH}} \equiv \tilde{w}^A$. Below, it will also be useful to note that we use $\tilde{h}_1 \equiv (\psi_{H_1}^0)^{\text{GH}}$ and $\tilde{h}_2 \equiv (\psi_{H_2}^0)^{\text{GH}}$ for the higgsinos.

2.4 Mixing Matrices

In the following, we describe in detail our conventions for neutralino, chargino, sfermion, and Higgs mixing. For purposes of cross-checking, we include in Appendix C expressions for the tree-level mass matrices of neutralinos, charginos, and third generation sfermions, as they appear in the set of conventions adopted here.

More importantly, essentially all SUSY spectrum calculators on the market today work with mass matrices which include higher-order corrections. Consequently, a formal dependence on the renormalization scheme and scale, and on the external momenta appearing in the corrections, enters the definition of the corresponding mixing matrices. Since, at the moment, no consensus exists on the most convenient definition to use here, the parameters appearing in blocks NMIX, UMI, VMIX, STOPMIX, SBOTMIX, STAUMIX, and ALPHA below should be thought of as ‘best choice’ solutions, at the discretion of each spectrum calculator. For example, one program may output on-shell parameters (with the external momenta e.g. corresponding to specific particle masses) in these blocks while another may be using $\overline{\text{DR}}$ definitions at certain ‘characteristic’ scales. For details on specific prescriptions, the manual of the particular spectrum calculator should be consulted.

Nonetheless, for obtaining loop-improved tree-level results, these parameters can normally be used as is. They can also be used for consistent cross section and decay width calculations at higher orders, but then the renormalization prescription employed by the spectrum calculator must match or be consistently matched to that of the intended higher order calculation.

Finally, different spectrum calculators may disagree on the overall sign of one or more rows in a mixing matrix, owing to different diagonalization algorithms. Such differences correspond to a flip of the sign of the eigenvectors in question and do not lead to inconsistencies. Only the relative sign between entries on the same row is physically significant, for processes with interfering amplitudes.

Neutralino Mixing

The Lagrangian contains the (symmetric) neutralino mass matrix as

$$\mathcal{L}_{\tilde{\chi}^0}^{\text{mass}} = -\frac{1}{2}\tilde{\psi}^{0T}\mathcal{M}_{\tilde{\psi}^0}\tilde{\psi}^0 + \text{h.c.} , \quad (10)$$

in the basis of 2-component spinors $\tilde{\psi}^0 = (-i\tilde{b}, -i\tilde{w}^3, \tilde{h}_1, \tilde{h}_2)^T$. We define the unitary 4 by 4 neutralino mixing matrix N (block NMIX below), such that:

$$-\frac{1}{2}\tilde{\psi}^{0T}\mathcal{M}_{\tilde{\psi}^0}\tilde{\psi}^0 = -\frac{1}{2}\underbrace{\tilde{\psi}^{0T}N^T}_{\tilde{\chi}^{0T}}\underbrace{N^*\mathcal{M}_{\tilde{\psi}^0}N^\dagger}_{\text{diag}(m_{\tilde{\chi}^0})}\underbrace{N\tilde{\psi}^0}_{\tilde{\chi}^0} , \quad (11)$$

where the (2-component) neutralinos $\tilde{\chi}_i^0$ are defined such that their absolute masses increase with increasing i . Generically, the resulting mixing matrix N may yield complex entries in the mass matrix, $\text{diag}(m_{\tilde{\chi}^0})_i = m_{\tilde{\chi}_i^0}e^{i\varphi_i}$. If so, we absorb the phase into the definition of the corresponding eigenvector, $\tilde{\chi}_i^0 \rightarrow \tilde{\chi}_i^0 e^{i\varphi_i/2}$, making the mass matrix strictly real:

$$\text{diag}(m_{\tilde{\chi}^0}) \equiv \left[N^*\mathcal{M}_{\tilde{\psi}^0}N^\dagger \right]_{ij} = m_{\tilde{\chi}_i^0}\delta_{ij} . \quad (12)$$

Note, however, that a special case occurs when CP violation is absent and one or more of the $m_{\tilde{\chi}_i^0}$ turn out to be negative. In this case, we allow for maintaining a strictly real mixing matrix N , instead writing the *signed* mass eigenvalues in the output. Thus, a negative $m_{\tilde{\chi}_i^0}$ in the output implies that the physical field is obtained by the rotation $\tilde{\chi}_i^0 \rightarrow \tilde{\chi}_i^0 e^{i\pi_i/2}$.

Our conventions on this point are slightly different from those used in [21] where the same Lagrangian appears but where the diagonalizing matrix $(N)^{\text{GH}}$ is chosen such that the elements of the diagonal mass matrix, $(N_D)^{\text{GH}}$ (the mass matrix, in the notation of [21]), are *always* real and non-negative, i.e. $(N_D)^{\text{GH}} \equiv |\text{diag}(m_{\tilde{\chi}^0})|$, at the price of $(N)^{\text{GH}}$ generally being complex-valued also in the absence of CP violation.

Chargino Mixing

We make the identification $\tilde{w}^\pm = (\tilde{w}^1 \mp i\tilde{w}^2)/\sqrt{2}$ for the charged winos and $\tilde{h}_1^-, \tilde{h}_2^+$ for the charged higgsinos. The Lagrangian contains the chargino mass matrix as

$$\mathcal{L}_{\tilde{\chi}^\pm}^{\text{mass}} = -\frac{1}{2}\tilde{\psi}^{-T}\mathcal{M}_{\tilde{\psi}^\pm}\tilde{\psi}^\pm + \text{h.c.} , \quad (13)$$

in the basis of 2-component spinors $\tilde{\psi}^+ = (-i\tilde{w}^+, \tilde{h}_2^+)^T$, $\tilde{\psi}^- = (-i\tilde{w}^-, \tilde{h}_1^-)^T$. We define the unitary 2 by 2 chargino mixing matrices, U and V (blocks UMIX

and VMIX below), such that:

$$-\frac{1}{2}\tilde{\psi}^{-T}\mathcal{M}_{\tilde{\psi}^+}\tilde{\psi}^+ = -\frac{1}{2}\underbrace{\tilde{\psi}^{-T}U^T}_{\tilde{\chi}^{-T}}\underbrace{U^*\mathcal{M}_{\tilde{\psi}^+}V^\dagger}_{\text{diag}(m_{\tilde{\chi}^+})}\underbrace{V\tilde{\psi}^+}_{\tilde{\chi}^+}, \quad (14)$$

where the (2-component) charginos $\tilde{\chi}_i^\pm$ are defined such that their absolute masses increase with increasing i and such that the mass matrix, $m_{\tilde{\chi}_i^\pm}$, is strictly real:

$$\text{diag}(m_{\tilde{\chi}^+}) \equiv \left[U\mathcal{M}_{\tilde{\psi}^+}V^T \right]_{ij} = m_{\tilde{\chi}_i^+}\delta_{ij}. \quad (15)$$

Again, in the absence of CP violation U and V can be chosen strictly real. This choice, as compared to that adopted in [21], shows similar differences as for neutralino mixing.

Sfermion Mixing

At present, we restrict our attention to left–right mixing in the third generation sfermion sector only. The convention we use is, for the interaction eigenstates, that \tilde{f}_L and \tilde{f}_R refer to the $SU(2)_L$ doublet and singlet superpartners of the fermion $f \in \{t, b, \tau\}$, respectively, and, for the mass eigenstates, that \tilde{f}_1 and \tilde{f}_2 refer to the lighter and heavier mass eigenstates, respectively. With this choice of basis, the spectrum output (blocks STOPMIX, SBOTMIX, and STAUMIX below) should contain the elements of the following matrix:

$$\begin{pmatrix} \tilde{f}_1 \\ \tilde{f}_2 \end{pmatrix} = \begin{bmatrix} F_{11} & F_{12} \\ F_{21} & F_{22} \end{bmatrix} \begin{pmatrix} \tilde{f}_L \\ \tilde{f}_R \end{pmatrix}, \quad (16)$$

whose determinant should be ± 1 . We here deliberately avoid notation involving mixing angles, to prevent misunderstandings which could arise due to the different conventions for these angles used in the literature. The mixing matrix elements themselves are unambiguous, apart from the overall signs of rows in the matrices, see above. Note that in [21], the mass eigenstates are *not* necessarily ordered in mass.

Higgs Mixing

The conventions for μ , v_1 , v_2 , v , $\tan\beta$, and m_A^2 were defined above in sections 2.2 and 2.3. The angle α (block ALPHA) we define by the rotation matrix:

$$\begin{pmatrix} H^0 \\ h^0 \end{pmatrix} = \begin{bmatrix} \cos\alpha & \sin\alpha \\ -\sin\alpha & \cos\alpha \end{bmatrix} \begin{pmatrix} H_1^0 \\ H_2^0 \end{pmatrix}, \quad (17)$$

where H_1^0 and H_2^0 are the CP–even neutral Higgs scalar interaction eigenstates, and h^0 and H^0 the corresponding mass eigenstates (including any higher order corrections present in the spectrum calculation), with $m_{h^0} < m_{H^0}$ by definition. This convention is identical to that of [21].

2.5 Running Couplings

In contrast to the effective definitions adopted above for the mixing matrices, we choose to define the parameters which appear in the output blocks `HMIX`, `GAUGE`, `M_SOFT`, `AU`, `AD`, `AE_YU`, `YD`, and `YE`, as $\overline{\text{DR}}$ running parameters, computed at one or more user-specifiable scales Q_i .

That the $\overline{\text{DR}}$ scheme is adopted for the output of running parameters is simply due to the fact that this scheme substantially simplifies many SUSY calculations (and hence all spectrum calculators use it). However, it does have drawbacks which for some applications are serious. For example, the $\overline{\text{DR}}$ scheme violates mass factorization as used in QCD calculations [28]. For consistent calculation beyond tree-level of processes relying on this factorization, e.g. cross sections at hadron colliders, the $\overline{\text{MS}}$ scheme is the only reasonable choice. At the present level of calculational precision, this is fortunately not an obstacle, since at one loop, a set of parameters calculated in either of the two schemes can be consistently translated into the other [29]. Explicit prescriptions for how to do this are included in appendix B.

Note, however, that different spectrum calculators use different choices for the underlying particle content of the effective theory. The programs `SOFTSUSY` (v. 1.8), `SPHENO` (v. 2.1), and `SUSPECT` (v. 2.2) use the full MSSM spectrum at all scales, whereas in `ISAJET` (v. 7.69) an alternative prescription is followed, with different particles integrated out of the effective theory at different scales. Whatever the case, these couplings should *not* be used ‘as is’ in calculations performed in another renormalization scheme or where a different effective field content is assumed. Thus, unfortunately, ensuring consistency of the field content assumed in the effective theory must still be done on a per program basis, though information on the prescription used by a particular spectrum calculator may conveniently be given in block `SPINFO`, when running parameters are provided.

Technically, we treat running parameters in the output in the following manner: since programs outside the spectrum calculation will not normally be able to run parameters with the full spectrum included, or at least less precisely than the spectrum calculators themselves, an option is included in block `MODSEL` below to instruct the spectrum calculator to write out values for each running parameter at a user-defined number of logarithmically spaced scales, i.e. to give output on running parameters at a grid of scales, Q_i , where the lowest point in the grid will normally be $Q_{\min} = m_Z$ and the highest point is user-specifiable. A complementary possibility is to let the spectrum calculator give output for the running couplings at one or more scales equal to specific sparticle masses in the spectrum. This option is also invoked using block `MODSEL`.

Warning: please note that these options are merely intended to *allow* information on running parameters to be passed, if desired. Many of the codes involved will at present not actually make use of this information, even if provided with it.

3 Definitions of the Interfaces

In this section, the SUSY Les Houches Accord input and output files are described. We here concentrate on the technical structure only. The reader should consult section 2 for parameter definitions and convention choices.

The following general structure for the SLHA files is proposed:

- All quantities with dimensions of energy (mass) are implicitly understood to be in GeV (GeV/c^2).
- Particles are identified by their PDG particle codes. See appendix A for lists of these, relevant to the MSSM. For a complete listing, see [22, chp. 33].
- The first character of every line is reserved for control and comment statements. Data lines should have the first character empty.
- In general, formatted output should be used for write-out, to avoid ‘messy-looking’ files, while a free format should be used on read-in, to avoid misalignment etc. leading to program crashes.
- Read-in should be performed in a case-insensitive way, again to increase stability.
- The general format for all real numbers is the FORTRAN format E16.8³. This large number of digits is used to avoid any possible numerical precision issue, and since it is no more difficult for e.g. the spectrum calculator to write out such a number than a shorter version. For typed input, it merely means that at least 16 spaces are reserved for the number, but e.g. the number 123.456 may be typed in “as is”. See also the example files in appendix D.
- A “#” mark anywhere means that the rest of the line is intended as a comment to be ignored by the reading program.
- All input and output is divided into sections in the form of named ‘blocks’. A “BLOCK xxxx” (with the “B” being the first character on the line) marks the beginning of entries belonging to the block named “xxxx”. E.g. “BLOCK MASS” marks that all following lines until the next “BLOCK” (or “DECAY”) statement contain mass values, to be read in a specific format, intrinsic to the MASS block. The order of blocks is arbitrary, except that input blocks should always come before output blocks.
- Reading programs should skip over blocks that are not recognized, issuing a warning rather than crashing. Thereby, stability is increased and private

³E16.8: a 16-character wide real number in scientific notation, whereof 8 digits are decimals, e.g. “-0.12345678E+000”.

blocks can be constructed, for instance `BLOCK MYCODE` could contain some parameters that only the program `MYCODE` (or a special hack of it) needs, but which are not recognized universally.

- A line with a blank first character is a data statement, to be interpreted according to what data the current block contains. Comments and/or descriptions added after the data values, e.g. “... `# comment`”, should always be added, to increase readability of the file for human readers.
- Use of the ‘tab’ character is dangerous and should be avoided.

Finally, program authors are advised to check that any parameter relations they assume in their codes (implicit or explicit) are either obeyed by the parameters in the files or disabled. As a specific example, take a code that normally would use e.g. the tree-level expression for the stop mixing matrix to compute A_t , given the stop mixing angle (together with a given set of other input parameters). This relation should *not* be used when reading in an SLHA spectrum; there may be (higher-order) contributions included in the spectrum calculation which cannot be absorbed into redefinitions of the tree-level couplings. The reading program should in this case read both the mixing matrix *and* A_t directly from the spectrum file, without assuming any a priori relation between them.

3.1 The Model Input File

Here, the user sets up his or her calculation, with low-energy boundary conditions and SUSY model parameters. If some or all of the low-energy boundary conditions are not supplied, the spectrum calculator should use its own defaults for those parameters, passing them on to the output file, so that the complete set of parameters that has been used for the calculation is available in the spectrum output. If the spectrum calculator has hard-coded defaults which the user is not allowed to change, the parameters that were *actually used* for the run should be written onto the output file.

The following general structure for the model input file is proposed:

- `BLOCK MODSEL`: Program-independent model switches, e.g. which model of supersymmetry breaking to use.
- `BLOCK SMINPUTS`: Measured values of SM parameters, used as boundary conditions in the spectrum calculation. These are also required for subsequent calculations to be consistent with the spectrum calculation.
- `BLOCK MINPAR`: Input parameters for minimal/default models. No defaults are defined for this block, and so the user must supply all required parameters. If `MINPAR` is not present, then *all* model parameters must be specified explicitly using `EXTPAR` below.

- **BLOCK EXTPAR:** Optional input parameters for non-minimal/non-universal models. This block may be entirely absent from the input file, in which case a minimal type of the selected SUSY breaking model will be used.

See also the example model input file included in appendix D.1.

BLOCK MODSEL

Switches and options for model selection. The entries in this block should consist of an index, identifying the particular switch in the listing below, followed by another integer or real number, specifying the option or value chosen. Switches so far defined are:

- 1 : Choice of SUSY breaking model. By default, a minimal type of model will always be assumed. Possible values are:
 - 0 : General MSSM Simulation.
 - 1 : (m)SUGRA model.
 - 2 : (m)GMSB model.
 - 3 : (m)AMSB model.
 - 4 : ...

- 3 : (Default=0) Choice of particle content. This switch is only meant as an example and is not yet implemented in any actual code. Switches defined could be:
 - 0 : MSSM.
 - 1 : NMSSM.
 - 2 : ...

- 11 : (Default=1) Number of points for a logarithmically spaced grid in Q for which the user wants the spectrum calculator to give output for the running parameters.
 - 1 : one copy of each block containing running parameters will be output, at the scale specified by Q_{\max} below.
 - $n > 1$: n copies of each block containing running parameters will be output. The smallest scale for the grid is normally m_Z while the maximum scale is set by Q_{\max} below.

- 12 : (Default= M_{EWSB}) Q_{max} . The largest Q scale at which to give the running parameters if a grid of output scales for each running parameter block has been requested using the switch above. The default is to give only one copy of each running parameter block, at the scale taken by the spectrum calculator to perform electroweak symmetry breaking. This is often taken to be $Q = M_{\text{EWSB}} \equiv \sqrt{m_{\tilde{t}_1} m_{\tilde{t}_2}}$.
- 21 : PDG code for a particle. The running SUSY-breaking mass parameters will be printed out at the pole mass of that particle, in addition to their values at scales given by the grid specified above. Several different entries can be given.

BLOCK SMINPUTS

Measured SM parameters, used as boundary conditions for the spectrum calculation and passed to subsequent calculations for consistency. Note that some programs have hard-coded defaults for various of these parameters, hence only a subset may sometimes be available as free inputs. The parameters, as defined in section 2.1, are:

- 1 : $\alpha_{\text{em}}^{-1}(m_Z)^{\overline{\text{MS}}}$. Inverse electromagnetic coupling at the Z pole in the $\overline{\text{MS}}$ scheme (with 5 active flavours).
- 2 : G_F . Fermi constant (in units of GeV^{-2}).
- 3 : $\alpha_s(m_Z)^{\overline{\text{MS}}}$. Strong coupling at the Z pole in the $\overline{\text{MS}}$ scheme (with 5 active flavours).
- 4 : m_Z , pole mass.
- 5 : $m_b(m_b)^{\overline{\text{MS}}}$. b quark running mass in the $\overline{\text{MS}}$ scheme.
- 6 : m_t , pole mass.
- 7 : m_τ , pole mass.

Please note that $m_b^{\text{pole}} \neq m_b(m_b)^{\overline{\text{MS}}} \neq m_b(m_b)^{\overline{\text{DR}}}$, see discussions in section 2.1 and [23, 24].

BLOCK MINPAR

Input parameters for minimal/default SUSY models. The interpretation given to the contents of this block naturally depends on which type of SUSY breaking model has been selected in block `MODEL`. Below are listed how `MINPAR` should be filled for `mSUGRA`, `mGMSB`, and `mAMSB` models, and for a general `MSSM`

setup. All parameters are understood to be $\overline{\text{DR}}$ parameters given at the input scale, M_{input} , which by default is the unification scale inferred from coupling unification. Alternatively, M_{input} can be given explicitly in block `EXTPAR` below. The only exception is $\tan\beta$ which we define (cf. section 2.2) as a $\overline{\text{DR}}$ parameter at the scale m_Z .

If a non-minimal type of model is desired, these minimal parameter sets may still be used to form the basis for the spectrum calculation, see `EXTPAR` below for details on this.

mSUGRA models.

- 1 : m_0 . Common scalar mass.
- 2 : $m_{1/2}$. Common gaugino mass.
- 3 : $\tan\beta$. Ratio of Higgs vacuum expectation, see section 2.2.
- 4 : $\text{sign}(\mu)$. Sign of the bilinear Higgs term in the superpotential.
- 5 : A . Common trilinear coupling.

mGMSB models.

- 1 : Λ . Scale of soft SUSY breaking felt by the low-energy sector.
- 2 : M_{mess} . Overall messenger scale.
- 3 : $\tan\beta$. Ratio of Higgs vacuum expectation values, see section 2.2.
- 4 : $\text{sign}(\mu)$. Sign of the bilinear Higgs term in the superpotential.
- 5 : N_5 . Messenger index.
- 6 : c_{grav} . Gravitino mass factor.

mAMSB models.

- 1 : m_0 . Common scalar mass term.
- 2 : $m_{3/2}$. Gravitino mass.
- 3 : $\tan\beta$. Ratio of Higgs vacuum expectation values, see section 2.2.
- 4 : $\text{sign}(\mu)$. Sign of the bilinear Higgs term in the superpotential.

Other models.

- 3 : $\tan\beta$. Ratio of Higgs vacuum expectation values, see section 2.2.

No model-specific standards for inputs for models beyond mSUGRA, mGMSB, and mAMSB have yet been defined, apart from the non-universality options available for these models in `EXTPAR` below. However, as long as a code for an alternative SUSY-breaking model adheres to the *output* standards described in the next section, there should be no problems in using it with this interface, as long as $\tan\beta$ is still provided.

BLOCK `EXTPAR`

Optional input parameters for non-minimal/non-universal models. This block may be entirely absent from the input file, in which case a minimal type of the selected SUSY breaking model will be used. When block `EXTPAR` is present, the starting point is still a minimal model with parameters as given in `MINPAR` but with each value present in `EXTPAR` replacing the minimal model value of that parameter, as applicable. If `MINPAR` is not present, then *all* model parameters must be specified explicitly using `EXTPAR`. All scale-dependent parameters are understood to be given in the $\overline{\text{DR}}$ scheme.

- 0 : M_{input} . Input scale for SUGRA, AMSB, and general MSSM models. If absent, the GUT scale derived from gauge unification will be used as input scale. Note that this parameter has no effect in GMSB scenarios where the input scale by definition is identical to the messenger scale, M_{mess} . A special case is when $Q = M_{\text{EWSB}} \equiv \sqrt{m_{\tilde{t}_1} m_{\tilde{t}_2}}$ is desired as input scale, since this scale is not known beforehand. This choice can be invoked by giving the special value $M_{\text{input}} = -1$.

Gaugino Masses

- 1 : $M_1(M_{\text{input}})$. $U(1)_Y$ gaugino (Bino) mass.
 2 : $M_2(M_{\text{input}})$. $SU(2)_L$ gaugino (Wino) mass.
 3 : $M_3(M_{\text{input}})$. $SU(3)_C$ gaugino (gluino) mass.

Trilinear Couplings

- 11 : $A_t(M_{\text{input}})$. Top trilinear coupling.
 12 : $A_b(M_{\text{input}})$. Bottom trilinear coupling.

13 : $A_\tau(M_{\text{input}})$. Tau trilinear coupling.

Higgs Parameters

— Either of the parameter sets $(m_{H_1}^2, m_{H_u}^2)$ or (μ, m_A^2) may be given, but not both.

- 21 : $m_{H_1}^2(M_{\text{input}})$. Down type Higgs mass squared.
 22 : $m_{H_2}^2(M_{\text{input}})$. Up type Higgs mass squared.
 23 : $\mu(M_{\text{input}})$. μ parameter.
 24 : $m_A^2(M_{\text{input}})$. Tree-level pseudoscalar Higgs mass squared.
 25 : $\tan\beta(M_{\text{input}})$. If present, this value of $\tan\beta$ overrides the one in MINPAR, and the input scale is taken as M_{input} rather than m_Z .

Sfermion Masses

- 31 : $m_{\tilde{e}_L}(M_{\text{input}})$. Left 1stgen. scalar lepton mass.
 32 : $m_{\tilde{\mu}_L}(M_{\text{input}})$. Left 2ndgen. scalar lepton mass.
 33 : $m_{\tilde{\tau}_L}(M_{\text{input}})$. Left 3rdgen. scalar lepton mass.
 34 : $m_{\tilde{e}_R}(M_{\text{input}})$. Right scalar electron mass.
 35 : $m_{\tilde{\mu}_R}(M_{\text{input}})$. Right scalar muon mass.
 36 : $m_{\tilde{\tau}_R}(M_{\text{input}})$. Right scalar tau mass.
 41 : $m_{\tilde{q}_{1L}}(M_{\text{input}})$. Left 1stgen. scalar quark mass.
 42 : $m_{\tilde{q}_{2L}}(M_{\text{input}})$. Left 2ndgen. scalar quark mass.
 43 : $m_{\tilde{q}_{3L}}(M_{\text{input}})$. Left 3rdgen. scalar quark mass.
 44 : $m_{\tilde{u}_R}(M_{\text{input}})$. Right scalar up mass.
 45 : $m_{\tilde{c}_R}(M_{\text{input}})$. Right scalar charm mass.
 46 : $m_{\tilde{t}_R}(M_{\text{input}})$. Right scalar top mass.
 47 : $m_{\tilde{d}_R}(M_{\text{input}})$. Right scalar down mass.
 48 : $m_{\tilde{s}_R}(M_{\text{input}})$. Right scalar strange mass.
 49 : $m_{\tilde{b}_R}(M_{\text{input}})$. Right scalar bottom mass.

Other Extensions

- 51 : $N_{5,1}$ (GMSB only). $U(1)_Y$ messenger index.
- 52 : $N_{5,2}$ (GMSB only). $SU(2)_L$ messenger index.
- 53 : $N_{5,3}$ (GMSB only). $SU(3)_C$ messenger index.

3.2 The Spectrum File

For the MSSM mass and coupling spectrum, the following block names are defined, to be specified further below:

- BLOCK MASS: Mass spectrum (pole masses).
- BLOCK NMIX: Neutralino mixing matrix.
- BLOCK UMIK: Chargino U mixing matrix.
- BLOCK VMIX: Chargino V mixing matrix.
- BLOCK STOPMIX: Stop mixing matrix.
- BLOCK SBOTMIX: Sbottom mixing matrix.
- BLOCK STAUMIX: Stau mixing matrix.
- BLOCK ALPHA: Higgs mixing angle α .
- BLOCK HMIX Q= . . . : μ , $\tan \beta$, v , and m_A^2 at scale Q .
- BLOCK GAUGE Q= . . . : Gauge couplings at scale Q .
- BLOCK MSOFT Q= . . . : Soft SUSY breaking mass parameters at scale Q .
- BLOCK AU, AD, AE Q= . . . : Trilinear couplings at scale Q .
- BLOCK YU, YD, YE Q= . . . : Yukawa couplings at scale Q .
- BLOCK SPINFO: Information from the spectrum calculator.

Note that there should always be at least one empty character between the BLOCK statement and the block name. For running parameters, an arbitrary number of each block may be written, to provide parameters at a grid of scales Q_i (in the $\overline{\text{DR}}$ scheme). For these blocks, the Q= statement should have at least one empty character on both sides. See also the example spectrum file included in appendix D.2.

BLOCK MASS

Mass spectrum for sparticles and Higgs bosons, signed pole masses. The standard for each line in the block should correspond to the FORTRAN format

$$(1x, I9, 3x, 1P, E16.8, 0P, 3x, '#', 1x, A),$$

where the first 9-digit integer should be the PDG code of a particle and the double precision number its mass.

BLOCK NMIX, UMIK, VMIX, STOPMIX, SBTMIX, STAUMIX

Mixing matrices, real parts only (CP violation is not addressed by this Accord at the present stage). The standard should correspond to the FORTRAN format

$$(1x, I2, 1x, I2, 3x, 1P, E16.8, 0P, 3x, '#', 1x, A).$$

For a generic mixing matrix X , the first two integers in the format represent i and j in X_{ij} respectively, and the double precision number X_{ij} itself. Note that different spectrum calculators may produce different overall signs for rows in these matrices, since an overall sign of an eigenvector does not change physics (see section 2.4 above).

BLOCK ALPHA

This block only contains one entry, the Higgs mixing angle α (see definition in Section 2.4), written in the format

$$(9x, 1P, E16.8, 0P, 3x, '#', 1x, A).$$
BLOCK HMIX Q= . . .

$\overline{\text{DR}}$ Higgs parameters at the scale Q , cf. sections 2.2 and 2.3. The entries in this block should consist of an index, identifying the particular parameter in the listing below, followed by a double precision number, giving the parameter value. The corresponding FORTRAN format would be

$$(1x, I5, 3x, 1P, E16.8, 0P, 3x, '#', 1x, A).$$

So far, the following entries are defined:

- 1 : $\mu(Q)$.
- 2 : $\tan \beta(Q)$.
- 3 : $v(Q)$.
- 4 : $m_A^2(Q)$.

BLOCK GAUGE Q= ...

$\overline{\text{DR}}$ gauge couplings at the scale Q , cf. section 2.2. The entries in this block should consist of an index, identifying the parameter in the listing below, the format being equivalent to that of block HMIX above. The parameters are:

- 1 : $g'(Q)$.
- 2 : $g(Q)$.
- 3 : $g_3(Q)$.

BLOCK MSOFT Q= ...

$\overline{\text{DR}}$ soft SUSY breaking mass parameters at the scale Q , cf. eqs. (7) and (9). The entries in this block should consist of an index, identifying the parameter in the listing below, the format being equivalent to that of block HMIX above and the numbering paralleling that of block EXTPAR. The parameters defined are:

- 1 : $M_1(Q)$.
- 2 : $M_2(Q)$.
- 3 : $M_3(Q)$.
- 21 : $m_{H_1}^2(Q)$.
- 22 : $m_{H_2}^2(Q)$.
- 31 : $m_{\tilde{e}_L}(Q)$.
- 32 : $m_{\tilde{\mu}_L}(Q)$.
- 33 : $m_{\tilde{\tau}_L}(Q)$.
- 34 : $m_{\tilde{e}_R}(Q)$.
- 35 : $m_{\tilde{\mu}_R}(Q)$.
- 36 : $m_{\tilde{\tau}_R}(Q)$.
- 41 : $m_{\tilde{q}_{1L}}(Q)$.
- 42 : $m_{\tilde{q}_{2L}}(Q)$.
- 43 : $m_{\tilde{q}_{3L}}(Q)$.
- 44 : $m_{\tilde{u}_R}(Q)$.
- 45 : $m_{\tilde{c}_R}(Q)$.
- 46 : $m_{\tilde{t}_R}(Q)$.

47 : $m_{\bar{d}_R}(Q)$.

48 : $m_{\bar{s}_R}(Q)$.

49 : $m_{\bar{b}_R}(Q)$.

BLOCK AU, AD, AE Q= . . .

$\overline{\text{DR}}$ soft breaking trilinear couplings at the scale Q , cf. eq. (5). These blocks are indexed like matrices (formatted like block MMIX above). At present, only the (3,3) component of each of these blocks should be given, corresponding to A_t , A_b , and A_τ , respectively. Other non-zero components would in general introduce mixing in the first and second generations, a situation which cannot be handled by the present Accord. This possibility is, however, left open for future development and/or private extensions.

BLOCK YU, YD, YE Q= . . .

$\overline{\text{DR}}$ fermion Yukawa couplings at the scale Q , cf. eq. (3). These blocks are indexed like matrices (formatted like block MMIX above). At present, only the (3,3) component of each of these blocks should be given, corresponding to the top quark, bottom quark, and tau lepton Yukawa couplings, respectively. Comments similar to those for the trilinear couplings above apply.

BLOCK SPINFO

Information from spectrum calculator. The program name and version number are obligatory. Optional: warnings and error messages from spectrum calculation, status return codes, regularization and renormalization prescription etc. The format should be

(1x,I5,3x,A).

Entries so far defined are:

- 1 : spectrum calculator name (string).
- 2 : spectrum calculator version number (string).
- 3 : If this entry is present, warning(s) were produced by the spectrum calculator. The resulting spectrum may still be OK. The entry should contain a description of the problem (string).
- 4 : If this entry is present, error(s) were produced by the spectrum calculator. The resulting spectrum should not be used. The entry should contain a description of the problem (string).

To illustrate, a certain unlucky choice of input parameters could result in the following form of block SPINFO:

```

BLOCK SPINFO
  1   MyRGE
  2   1.0
  3   Charge and colour breaking global minimum
  3   Desired accuracy not quite achieved
  3   LEP2 Higgs bound violated
  4   No radiative electroweak symmetry breaking
  4   Tachyons encountered

```

3.3 The Decay File

The decay table for each particle begins with a statement specifying which particle is decaying and its total width, in the format:

```

#           PDG           Width
DECAY  1000021   1.01752300e+00 # gluino decays

```

The first integer is a PDG particle number, specifying the identity of the mother of all subsequent lines until the next DECAY or BLOCK statement (or end-of-file). The subsequent real number is that particle's total width. The end comment contains a human readable translation of the PDG code.

Every subsequent line contains a decay channel for this mother in the format:

```

#           BR           NDA           ID1           ID2
  4.18313300E-02   2   1000001   -1 # BR(~g -> ~d_L dbar)
  1.55587600E-02   2   2000001   -1 # BR(~g -> ~d_R dbar)
  3.91391000E-02   2   1000002   -2 # BR(~g -> ~u_L ubar)
  1.74358200E-02   2   2000002   -2 # BR(~g -> ~u_R ubar)
...

```

where the first real number is the fraction of the total width (branching fraction) into that particular mode, the first integer is the number of daughters, N , and the N following integers are the PDG codes of the N daughters. The specific FORTRAN formats for the DECAY statement and the entries in the decay table are, respectively:

```

('DECAY', 1x, I9, 3x, 1P, E16.8, 0P, 3x, '#', 1x, A),
(3x, 1P, E16.8, 0P, 3x, I2, 3x, N(I9, 1x), 2x, '#', 1x, A).

```

A potential pitfall in using these decay tables is how on-shell resonances inside the physical phase space are dealt with. A prime example, which we will use for illustration below, is the top decay, $t \rightarrow W^+ b \rightarrow q\bar{q}'b$. There are two dangers here that must be consistently dealt with:

1. If both the partial width for $t \rightarrow W^+b$ and for $t \rightarrow q\bar{q}'b$ are written on the file, then the on-shell W part of $t \rightarrow q\bar{q}'b$ will be counted twice. One solution here is to include only the truly off-shell parts in the calculation of partial widths for processes which can occur via sequences of (lower-order) on-shell splittings. In the example above, the $(1 \rightarrow 3)$ partial width $t \rightarrow q\bar{q}'b$ written on the file should thus *not* contain the $(1 \rightarrow 2 \otimes 1 \rightarrow 2)$ on-shell W contribution.

2. If the on-shell/off-shell rule just described is *not* adhered to, another possible source of error becomes apparent. If the full partial width $t \rightarrow q\bar{q}'b$ (now including the resonant piece) is written in the decay table, the reading program has no immediate way of knowing that there is a resonant W inside. By default, it will most likely use a flat phase space, in the lack of more differential information; hence there would in this example be no W peak in the $q\bar{q}'$ invariant mass spectrum. On the other hand, if the rule above *is* adhered to, then the resonant W is not part of the $1 \rightarrow 3$ partial width, and a flat phase space is then a reasonable first approximation.

NB: for Majorana particles, modes and charge conjugate modes should both be written on the file, so that the numbers in the first column sum up to 1 for any particle.

At present, this file is thus only capable of transferring integrated information, i.e. partial widths. For a more accurate population of phase space (even when all intermediate states are off-shell), access to differential information is necessary. It should also here be mentioned that several programs include options for letting the partial widths be a function of \hat{s} , to account for the resonant shape of the mother. We anticipate that these and other refinements can be included with full backwards compatibility, either by continuing to add information on the same line before the hash mark, or by adding a number of lines beginning with the '+' character for each decay mode in question, where additional information concerning that mode can be given.

The file should also contain a block DCINFO, giving information about the decay calculation program and (optionally) exit status. The format and entries of this block are identical to that of block SPINFO above.

Note that, if the particle is not known to the reading program, properties such as electric charge, spin, and colour charge could be added on the DECAY line. However, for fundamental particles, there is only a limited set of particles for which all programs concerned should know the properties, and for hadrons, the spin is encoded in the last digit of the PDG code while the charge may be calculated from the flavour content which is specified by the 2 or 3 digits preceding the spin digit of the code. Therefore, we do not deem it necessary to adopt a standard for specifying such information at the present time.

4 Conclusion

The present Accord specifies a unique set of conventions together with ASCII file formats for model input and spectrum output for most commonly investigated supersymmetric models, as well as a decay table file format for use with decay packages.

With respect to the model parameter input file, mSUGRA, mGMSB, and mAMSB scenarios can be handled, with some options for non-universality. However, this should not discourage users desiring to investigate alternative models; the definitions for the spectrum output file are at present capable of handling any CP and R-parity conserving supersymmetric model, with the particle spectrum of the MSSM. Specifically, this includes the so-called SPS points [30].

Also, these definitions are not intended to be static solutions. Great efforts have gone into ensuring that the Accord may accommodate essentially any new model or new twist on an old one with minor modifications required and full backwards compatibility. Planned issues for future extensions of the Accord are, for instance, to include options for R-parity violation and CP violation, and possibly to include definitions for an NMSSM. Topics which are at present only implemented in a few codes, if at all, will be taken up as the need arises. Handling RPV and CPV should require very minor modifications to the existing structure, while the NMSSM, for which there is at present not even general agreement on a unique definition, will require some additional work.

Lastly, while SUSY is perhaps the most studied hypothesis of New Physics, it is by no means the only possible worth investigating. One may well anticipate that similar sources of confusion and misunderstandings as partly motivated this Accord can arise for other New Physics models in the future. In this context, we note that the decay tables defined here are sufficiently general to require little or no modification to encompass other New Physics models, while the rest of the Accord provides a general structure that may be used as a convenient template for future generalisations.

Acknowledgements

The authors are grateful to the organizers of the Physics at TeV Colliders workshop (Les Houches, 2003) and to the organizers of the Workshop on Monte Carlo tools for the LHC (MC4LHC, CERN, 2003). The discussions and agreements reached at those two workshops constitute the backbone of this writeup.

This work has been supported in part by CERN, by the “Collider Physics” European Network under contract HPRN-CT-2000-00149, by the Swiss Bundesamt für Bildung und Wissenschaft, and by the National Science Foundation under Grant No. PHY99-07949. Research at ANL is supported by the U.S. Department of Energy HEP Division under the contract W-31-109-ENG-38.

W.P. is supported by the Erwin Schrödinger fellowship No. J2272 of the 'Fonds zur Förderung der wissenschaftlichen Forschung' of Austria and partly by the Swiss 'Nationalfonds'.

A The PDG Particle Numbering Scheme

Listed in the tables below are the PDG codes for the MSSM particle spectrum. Codes for other particles may be found in [22, chp. 33].

Table 1: SM fundamental particle codes (+ extended Higgs sector).

Code	Name	Code	Name	Code	Name	Code	Name
1	d	11	e^-	21	g		
2	u	12	ν_e	22	γ	35	H^0
3	s	13	μ^-	23	Z^0	36	A^0
4	c	14	ν_μ	24	W^+	37	H^+
5	b	15	τ^-	25	h^0		
6	t	16	ν_τ			39	G (graviton)

Table 2: MSSM sparticle codes

Code	Name	Code	Name	Code	Name	Code	Name
1000001	\tilde{d}_L	1000011	\tilde{e}_L	1000021	\tilde{g}		
1000002	\tilde{u}_L	1000012	$\tilde{\nu}_{eL}$	1000022	$\tilde{\chi}_1^0$	1000035	$\tilde{\chi}_4^0$
1000003	\tilde{s}_L	1000013	$\tilde{\mu}_L$	1000023	$\tilde{\chi}_2^0$		
1000004	\tilde{c}_L	1000014	$\tilde{\nu}_{\mu L}$	1000024	$\tilde{\chi}_1^+$	1000037	$\tilde{\chi}_2^+$
1000005	\tilde{b}_1	1000015	$\tilde{\tau}_1$	1000025	$\tilde{\chi}_3^0$		
1000006	\tilde{t}_1	1000016	$\tilde{\nu}_{\tau L}$			1000039	\tilde{G} (gravitino)
2000001	\tilde{d}_R	2000011	\tilde{e}_R				
2000002	\tilde{u}_R	2000012	$\tilde{\nu}_{eR}$				
2000003	\tilde{s}_R	2000013	$\tilde{\mu}_R$				
2000004	\tilde{c}_R	2000014	$\tilde{\nu}_{\mu R}$				
2000005	\tilde{b}_2	2000015	$\tilde{\tau}_2$				
2000006	\tilde{t}_2	2000016	$\tilde{\nu}_{\tau R}$				

B 1-Loop Translations from $\overline{\text{DR}}$ to $\overline{\text{MS}}$

All formulae in this appendix are obtained from [29] and are valid for parameters calculated at 1 loop in either of the two schemes.

At the scale μ , the $\overline{\text{MS}}$ gauge coupling of the gauge group G_i (i.e. the coupling appearing in the interaction vertices of the corresponding gauge bosons) is related to the $\overline{\text{DR}}$ one by:

$$g_{i,\overline{\text{MS}}} = g_{i,\overline{\text{DR}}} \left(1 - \frac{g_i^2}{96\pi^2} C(G_i) \right), \quad (18)$$

where the choice of renormalization scheme for g_i in the 1-loop correction piece is irrelevant, since it amounts to a 2-loop effect, and $C(G_i)$ is the quadratic Casimir invariant for the adjoint representation of the gauge group in question.⁴

For the running Yukawa couplings (between the scalar ϕ_i and the two chiral fermions ψ_j and ψ_k), the translation is:

$$Y_{\overline{\text{MS}}}^{ijk} = Y_{\overline{\text{DR}}}^{ijk} \left(1 + \sum_{a=1}^3 \frac{g_a^2}{32\pi^2} [C_a(r_j) - 2C_a(r_i) + C_a(r_k)] \right), \quad (19)$$

where a runs over the SM gauge groups. Note that while the $\overline{\text{DR}}$ Yukawas are totally symmetric, this is not the case for the $\overline{\text{MS}}$ ones. The same relation can also be used to derive the translation of fermion masses coming from a quadratic term in the superpotential by taking the scalar field as a dummy field with $C(r_i) = 0$ and identifying $C(r_j) = C(r_k)$. This applies e.g. to the Higgs mixing parameter μ .

The soft supersymmetry-breaking parameters differ in the two schemes, too. The relation between the gaugino masses is given by

$$M_{i,\overline{\text{MS}}} = M_{i,\overline{\text{DR}}} \left(1 + \frac{g_i^2}{16\pi^2} C(G_i) \right), \quad (20)$$

Finally, none of the other supersymmetry-breaking couplings (as written in component rather than superfield notation) differ between the two schemes. In particular, this applies to the soft breaking trilinear couplings and the scalar masses, provided one uses the modified $\overline{\text{DR}}$ scheme, as presented in [27].

⁴It should be noted that the couplings of gauginos to scalars are identical to the gauge couplings by virtue of supersymmetry. This requires the introduction of additional counter terms in the $\overline{\text{MS}}$ scheme [29] in order to restore this equality. Analogous additional counter terms arise for the quartic scalar couplings which are related to the Yukawa and gauge couplings due to supersymmetry.

C Tree-level Mass Matrices

The following gives a list of tree-level mass matrices, as they appear in the conventions adopted in this article, see Section 2. Note that these expressions *are not* normally the ones used for actual calculations in the spectrum calculators, since most codes on the market today include higher order corrections which are absent below.

The neutralino mass matrix appearing in eq. (11) is:

$$\mathcal{M}_{\tilde{\psi}0} = \begin{pmatrix} M_1 & 0 & -m_Z \cos \beta \sin \theta_W & m_Z \sin \beta \sin \theta_W \\ 0 & M_2 & m_Z \cos \beta \cos \theta_W & -m_Z \sin \beta \cos \theta_W \\ -m_Z \cos \beta \sin \theta_W & m_Z \cos \beta \cos \theta_W & 0 & -\mu \\ m_Z \sin \beta \sin \theta_W & -m_Z \sin \beta \cos \theta_W & -\mu & 0 \end{pmatrix}, \quad (21)$$

and the chargino mass matrix appearing in eq. (14):

$$\mathcal{M}_{\tilde{\psi}\pm} = \begin{pmatrix} M_2 & \sqrt{2}m_W \sin \beta \\ \sqrt{2}m_W \cos \beta & \mu \end{pmatrix}. \quad (22)$$

For the sfermions, the mixing matrices for \tilde{t} , \tilde{b} , and $\tilde{\tau}$ respectively, appear in the L-R basis as:

$$\begin{pmatrix} m_{\tilde{q}_{3L}}^2 + m_t^2 + \left(\frac{1}{2} - \frac{2}{3} \sin^2 \theta_W\right) m_Z^2 \cos 2\beta & m_t (A_t - \mu \cot \beta) \\ m_t (A_t - \mu \cot \beta) & m_{\tilde{t}_R}^2 + m_t^2 + \frac{2}{3} \sin^2 \theta_W m_Z^2 \cos 2\beta \end{pmatrix}, \quad (23)$$

$$\begin{pmatrix} m_{\tilde{q}_{3L}}^2 + m_b^2 - \left(\frac{1}{2} - \frac{1}{3} \sin^2 \theta_W\right) m_Z^2 \cos 2\beta & m_b (A_b - \mu \tan \beta) \\ m_b (A_b - \mu \tan \beta) & m_{\tilde{b}_R}^2 + m_b^2 - \frac{1}{3} \sin^2 \theta_W m_Z^2 \cos 2\beta \end{pmatrix}, \quad (24)$$

$$\begin{pmatrix} m_{\tilde{\tau}_L}^2 + m_\tau^2 - \left(\frac{1}{2} - \sin^2 \theta_W\right) m_Z^2 \cos 2\beta & m_\tau (A_\tau - \mu \tan \beta) \\ m_\tau (A_\tau - \mu \tan \beta) & m_{\tilde{\tau}_R}^2 + m_\tau^2 - \sin^2 \theta_W m_Z^2 \cos 2\beta \end{pmatrix}, \quad (25)$$

where we use $m_{\tilde{q}_{3L}}$ to denote the 3rd generation left squark mass.

D Examples

D.1 Example Input File

In the example below, the user has not entered boundary values for the electroweak couplings, nor have the Z and τ masses been supplied. On running, the spectrum calculator should thus use its own defaults for these parameters and pass everything on to the output.

```
# SUSY Les Houches Accord 1.0 - example input file
# Snowsas point 1a
Block MODSEL # Select model
  1 1 # sugra
Block SMINPUTS # Standard Model inputs
  3 0.1172 # alpha_s(MZ) SM MSbar
  5 4.25 # Mb(mb) SM MSbar
  6 174.3 # Mtop(pole)
Block MINPAR # SUSY breaking input parameters
  3 10.0 # tanb
  4 1.0 # sign(mu)
  1 100.0 # m0
  2 250.0 # m12
  5 -100.0 # A0
```

D.2 Example Spectrum File

The spectrum file produced by the above input file should look something like the following:

```
# SUSY Les Houches Accord 1.0 - example spectrum file
# Info from spectrum calculator
Block SPINFO # Program information
  1 SOFTSUSY # spectrum calculator
  2 1.8.4 # version number
# Input parameters
Block MODSEL # Select model
  1 1 # sugra
Block SMINPUTS # Standard Model inputs
  1 1.27934000e+02 # alpha_em^(-1)(MZ) SM MSbar
  2 1.16637000e-05 # G_Fermi
  3 1.17200000e-01 # alpha_s(MZ) SM MSbar
  4 9.11876000e+01 # MZ(pole)
```

```

5 4.25000000e+00 # Mb(mb) SM MSbar
6 1.74300000e+02 # Mtop(pole)
7 1.77700000e+00 # Mtau(pole)
Block MINPAR # SUSY breaking input parameters
1 1.00000000e+02 # m0(MGUT) MSSM DRbar
2 2.50000000e+02 # m12(MGUT) MSSM DRbar
3 1.00000000e+01 # tanb(MZ) MSSM DRbar
4 1.00000000e+00 # sign(mu(MGUT)) MSSM DRbar
5 -1.00000000e+02 # A0(MGUT) MSSM DRbar
#
# mgut=2.551299875e+16 GeV
Block MASS # Mass spectrum
# PDG code mass particle
24 8.02463984e+01 # MW
25 1.10636832e+02 # h0
35 4.00874604e+02 # H0
36 4.00506272e+02 # A0
37 4.08784776e+02 # H+
1000001 5.73103437e+02 # ~d_L
1000002 5.67658152e+02 # ~u_L
1000003 5.73029886e+02 # ~s_L
1000004 5.67583798e+02 # ~c_L
1000005 5.15617364e+02 # ~b_1
1000006 3.96457239e+02 # ~t_1
1000011 2.04346872e+02 # ~e_L
1000012 1.88733599e+02 # ~nue_L
1000013 2.04344144e+02 # ~mu_L
1000014 1.88730645e+02 # ~numu_L
1000015 1.36434250e+02 # ~stau_1
1000016 1.87868618e+02 # ~nu_tau_L
1000021 6.09298476e+02 # ~g
1000022 9.62723703e+01 # ~neutralino(1)
1000023 1.79383645e+02 # ~neutralino(2)
1000024 1.78832499e+02 # ~chargino(1)
1000025 -3.64094094e+02 # ~neutralino(3)
1000035 3.82627159e+02 # ~neutralino(4)
1000037 3.82906800e+02 # ~chargino(2)
2000001 5.46469067e+02 # ~d_R
2000002 5.47396607e+02 # ~u_R
2000003 5.46466786e+02 # ~s_R
2000004 5.47244191e+02 # ~c_R
2000005 5.44364792e+02 # ~b_2
2000006 5.85965936e+02 # ~t_2
2000011 1.45533727e+02 # ~e_R
2000013 1.45526066e+02 # ~mu_R
2000015 2.08212282e+02 # ~stau_2
# Higgs mixing

```

```

Block alpha # Effective Higgs mixing parameter
          -1.13716828e-01 # alpha
Block stopmix # stop mixing matrix
  1  1    5.37975095e-01 # O_{11}
  1  2    8.42960733e-01 # O_{12}
  2  1    8.42960733e-01 # O_{21}
  2  2   -5.37975095e-01 # O_{22}
Block sbotmix # sbottom mixing matrix
  1  1    9.47346882e-01 # O_{11}
  1  2    3.20209128e-01 # O_{12}
  2  1   -3.20209128e-01 # O_{21}
  2  2    9.47346882e-01 # O_{22}
Block stauxmix # stau mixing matrix
  1  1    2.78399839e-01 # O_{11}
  1  2    9.60465267e-01 # O_{12}
  2  1    9.60465267e-01 # O_{21}
  2  2   -2.78399839e-01 # O_{22}
Block nmix # neutralino mixing matrix
  1  1    9.86102610e-01 # N_{1,1}
  1  2   -5.46971979e-02 # N_{1,2}
  1  3    1.47526998e-01 # N_{1,3}
  1  4   -5.33445802e-02 # N_{1,4}
  2  1    1.01818619e-01 # N_{2,1}
  2  2    9.43310250e-01 # N_{2,2}
  2  3   -2.73948058e-01 # N_{2,3}
  2  4    1.57325147e-01 # N_{2,4}
  3  1   -6.06211640e-02 # N_{3,1}
  3  2    9.00367885e-02 # N_{3,2}
  3  3    6.95440071e-01 # N_{3,3}
  3  4    7.10339045e-01 # N_{3,4}
  4  1   -1.16446066e-01 # N_{4,1}
  4  2    3.14749686e-01 # N_{4,2}
  4  3    6.47727839e-01 # N_{4,3}
  4  4   -6.83974850e-01 # N_{4,4}
Block Umix # chargino U mixing matrix
  1  1    9.16207706e-01 # U_{1,1}
  1  2   -4.00703680e-01 # U_{1,2}
  2  1    4.00703680e-01 # U_{2,1}
  2  2    9.16207706e-01 # U_{2,2}
Block Vmix # chargino V mixing matrix
  1  1    9.72887524e-01 # V_{1,1}
  1  2   -2.31278762e-01 # V_{1,2}
  2  1    2.31278762e-01 # V_{2,1}
  2  2    9.72887524e-01 # V_{2,2}
Block gauge Q= 4.64649125e+02
  1    3.60872342e-01 # g'(Q)MSSM DRbar
  2    6.46479280e-01 # g(Q)MSSM DRbar

```

```

    3      1.09623002e+00 # g3(Q)MSSM DRbar
Block yu Q= 4.64649125e+02
    3 3      8.88194465e-01 # Yt(Q)MSSM DRbar
Block yd Q= 4.64649125e+02
    3 3      1.40135884e-01 # Yb(Q)MSSM DRbar
Block ye Q= 4.64649125e+02
    3 3      9.97405356e-02 # Ytau(Q)MSSM DRbar
Block hmix Q= 4.64649125e+02 # Higgs mixing parameters
    1      3.58660361e+02 # mu(Q)MSSM DRbar
    2      9.75139550e+00 # tan beta(Q)MSSM DRbar
    3      2.44923506e+02 # higgs vev(Q)MSSM DRbar
    4      1.69697051e+04 # mA^2(Q)MSSM DRbar
Block msoft Q=4.64649125e+02 # MSSM DRbar SUSY breaking pars
    1      1.01353084e+02 # M_1(Q)
    2      1.91513233e+02 # M_2(Q)
    3      5.86951218e+02 # M_3(Q)
   21      3.26601234e+04 # mH1^2(Q)
   22     -1.29761234e+05 # mH2^2(Q)
   31      1.99111011e+02 # meL(Q)
   32      1.99108212e+02 # mmuL(Q)
   33      1.98291304e+02 # mtauL(Q)
   34      1.38808102e+02 # meR(Q)
   35      1.38800070e+02 # mmuR(Q)
   36      1.36441129e+02 # mtauR(Q)
   41      5.51249170e+02 # mQL1(Q)
   42      5.51173571e+02 # mQL2(Q)
   43      4.99839787e+02 # mQL3(Q)
   44      5.29285249e+02 # muR(Q)
   45      5.29130286e+02 # mcR(Q)
   46      4.19025924e+02 # mtR(Q)
   47      5.26529352e+02 # mdR(Q)
   48      5.26527025e+02 # msR(Q)
   49      5.23183913e+02 # mbR(Q)
Block au Q= 4.64649125e+02
    3 3     -5.04995511e+02 # At(Q)MSSM DRbar
Block ad Q= 4.64649125e+02
    3 3     -7.97992485e+02 # Ab(Q)MSSM DRbar
Block ae Q= 4.64649125e+02
    3 3     -2.56328558e+02 # Atau(Q)MSSM DRbar

```

D.3 Example Decay File

For brevity, the model input and spectrum information is omitted here. See the examples above.

```

# SUSY Les Houches Accord 1.0 - example decay file
# Info from decay package
Block DCINFO          # Program information
  1   SDECAY          # Decay package
  2   1.0             # version number
#   PDG              Width
DECAY  1000021      1.01752300e+00 # gluino decays
#      BR           NDA      ID1      ID2
  4.18313300E-02   2      1000001     -1 # BR( $\tilde{g} \rightarrow \tilde{d}_L \text{ dbar}$ )
  1.55587600E-02   2      2000001     -1 # BR( $\tilde{g} \rightarrow \tilde{d}_R \text{ dbar}$ )
  3.91391000E-02   2      1000002     -2 # BR( $\tilde{g} \rightarrow \tilde{u}_L \text{ ubar}$ )
  1.74358200E-02   2      2000002     -2 # BR( $\tilde{g} \rightarrow \tilde{u}_R \text{ ubar}$ )
  4.18313300E-02   2      1000003     -3 # BR( $\tilde{g} \rightarrow \tilde{s}_L \text{ sbar}$ )
  1.55587600E-02   2      2000003     -3 # BR( $\tilde{g} \rightarrow \tilde{s}_R \text{ sbar}$ )
  3.91391000E-02   2      1000004     -4 # BR( $\tilde{g} \rightarrow \tilde{c}_L \text{ cbar}$ )
  1.74358200E-02   2      2000004     -4 # BR( $\tilde{g} \rightarrow \tilde{c}_R \text{ cbar}$ )
  1.13021900E-01   2      1000005     -5 # BR( $\tilde{g} \rightarrow \tilde{b}_1 \text{ bbar}$ )
  6.30339800E-02   2      2000005     -5 # BR( $\tilde{g} \rightarrow \tilde{b}_2 \text{ bbar}$ )
  9.60140900E-02   2      1000006     -6 # BR( $\tilde{g} \rightarrow \tilde{t}_1 \text{ tbar}$ )
  0.00000000E+00   2      2000006     -6 # BR( $\tilde{g} \rightarrow \tilde{t}_2 \text{ tbar}$ )
  4.18313300E-02   2      -1000001      1 # BR( $\tilde{g} \rightarrow \tilde{\text{dbar}}_L \text{ d}$ )
  1.55587600E-02   2      -2000001      1 # BR( $\tilde{g} \rightarrow \tilde{\text{dbar}}_R \text{ d}$ )
  3.91391000E-02   2      -1000002      2 # BR( $\tilde{g} \rightarrow \tilde{\text{ubar}}_L \text{ u}$ )
  1.74358200E-02   2      -2000002      2 # BR( $\tilde{g} \rightarrow \tilde{\text{ubar}}_R \text{ u}$ )
  4.18313300E-02   2      -1000003      3 # BR( $\tilde{g} \rightarrow \tilde{\text{sbar}}_L \text{ s}$ )
  1.55587600E-02   2      -2000003      3 # BR( $\tilde{g} \rightarrow \tilde{\text{sbar}}_R \text{ s}$ )
  3.91391000E-02   2      -1000004      4 # BR( $\tilde{g} \rightarrow \tilde{\text{cbar}}_L \text{ c}$ )
  1.74358200E-02   2      -2000004      4 # BR( $\tilde{g} \rightarrow \tilde{\text{cbar}}_R \text{ c}$ )
  1.13021900E-01   2      -1000005      5 # BR( $\tilde{g} \rightarrow \tilde{\text{bbar}}_1 \text{ b}$ )
  6.30339800E-02   2      -2000005      5 # BR( $\tilde{g} \rightarrow \tilde{\text{bbar}}_2 \text{ b}$ )
  9.60140900E-02   2      -1000006      6 # BR( $\tilde{g} \rightarrow \tilde{\text{tbar}}_1 \text{ t}$ )
  0.00000000E+00   2      -2000006      6 # BR( $\tilde{g} \rightarrow \tilde{\text{tbar}}_2 \text{ t}$ )

```

References

- [1] B. C. Allanach, *Comput. Phys. Commun.* **143**, 305 (2002).
- [2] H. Baer, F. E. Paige, S. D. Protopopescu, X. Tata, (1993), [[hep-ph/9305342](#)].

- [3] A. Djouadi, J.-L. Kneur, G. Moultaka, (2002), [hep-ph/0211331].
- [4] S. Heinemeyer, W. Hollik, G. Weiglein, *Comput. Phys. Commun.* **124**, 76 (2000),
(see www.feynhiggs.de).
- [5] W. Porod, *Comput. Phys. Commun.* **153**, 275 (2003).
- [6] W. Beenakker, R. Hopker, M. Spira, (1996), [hep-ph/9611232].
- [7] A. Djouadi, J. Kalinowski, M. Spira, *Comput. Phys. Commun.* **108**, 56 (1998).
- [8] M. Muhlleitner, A. Djouadi, Y. Mambrini, (2003), [hep-ph/0311167].
- [9] G. Belanger, F. Boudjema, A. Pukhov, A. Semenov, *Comput. Phys. Commun.* **149**, 103 (2002).
- [10] P. Gondolo et al., (2002), [astro-ph/0211238].
- [11] H. Baer, F. E. Paige, S. D. Protopopescu, X. Tata, (1999), [hep-ph/0001086].
- [12] G. Corcella et al., *JHEP* **01**, 010 (2001).
- [13] T. Gleisberg et al., (2003), [hep-ph/0311263].
- [14] A. Pukhov et al., (1999), [hep-ph/9908288].
- [15] T. Sjostrand et al., *Comput. Phys. Commun.* **135**, 238 (2001).
- [16] T. Sjostrand, L. Lonnblad, S. Mrenna, P. Skands, (2003), [hep-ph/0308153].
- [17] Minami-Tateya collaboration, H. Tanaka, M. Kuroda, T. Kaneko, M. Jimbo, T. Kon, *Nucl. Instrum. Meth.* **A389**, 295 (1997).
- [18] W. Kilian, LC-TOOL-2001-039.
- [19] CERN Computer News Letters (July-September 1994).
- [20] E. Boos et al., (2001), [hep-ph/0109068].
- [21] J. F. Gunion, H. E. Haber, *Nucl. Phys.* **B272**, 1 (1986), [Erratum-ibid. **B402**, 567 (1993)].
- [22] Particle Data Group, K. Hagiwara et al., *Phys. Rev.* **D66**, 010001 (2002).
- [23] K. Melnikov, T. v. Ritbergen, *Phys. Lett.* **B482**, 99 (2000).

-
- [24] H. Baer, J. Ferrandis, K. Melnikov, X. Tata, Phys. Rev. **D66**, 074007 (2002).
- [25] W. Siegel, Phys. Lett. **B84**, 193 (1979).
- [26] D. M. Capper, D. R. T. Jones, P. van Nieuwenhuizen, Nucl. Phys. **B167**, 479 (1980).
- [27] I. Jack, D. R. T. Jones, S. P. Martin, M. T. Vaughn, Y. Yamada, Phys. Rev. **D50**, 5481 (1994).
- [28] W. Beenakker, H. Kuijf, W. L. van Neerven, J. Smith, Phys. Rev. **D40**, 54 (1989).
- [29] S. P. Martin, M. T. Vaughn, Phys. Lett. **B318**, 331 (1993).
- [30] B. C. Allanach et al., Eur. Phys. J. **C25**, 113 (2002),
(see www.ipp.dur.ac.uk/~georg/sps/).

Measuring Neutrino Mixing angles at LHC

Paper IV

LU TP 03–50
December 2003

MEASURING NEUTRINO MIXING ANGLES AT LHC

W. Porod¹ and P. Skands²

¹Institut für Theoretische Physik, University of Zürich, CH-8057 Zürich,
Switzerland

²Dept. of Theoretical Physics, Lund University, SE-223 62 Lund, Sweden

Abstract We study an MSSM model with bilinear R-parity violation which is capable of explaining neutrino data while leading to testable predictions for ratios of LSP decay rates. Further, we estimate the precision with which such a measurement could be carried out at the LHC.

1 Introduction

Recent neutrino experiments [1–4] clearly show that neutrinos are massive particles and that they mix. In supersymmetric models these findings can be explained by the usual seesaw mechanism [5–7]. However, supersymmetry allows for an alternative which is intrinsically supersymmetric, namely the breaking of R-parity. The simplest way to realize this idea is to add bilinear terms to the superpotential W :

$$W = W_{\text{MSSM}} + \epsilon_i \hat{L}_i \hat{H}_u \quad (1)$$

For consistency one has also to add the corresponding bilinear terms to soft SUSY breaking which induce small vacuum expectation values (vevs) for the sneutrinos. These vevs in turn induce a mixing between neutrinos and neutralinos, giving mass to one neutrino at tree level. The second neutrino mass is induced by loop effects (see [8–10] and references therein). The same parameters that induce neutrino masses and mixings are also responsible for the decay of the lightest supersymmetric particle (LSP). This implies that there are correlations between neutrino physics and LSP decays [11–13].

In this note we investigate how well LHC can measure ratios of LSP branching ratios that are correlated to neutrino mixing angles in a scenario where the lightest neutralino $\tilde{\chi}_1^0$ is the LSP. In particular we focus on the semi-leptonic final states $l_i q' \bar{q}$ ($l_i = e, \mu, \tau$). There are several more examples which are discussed in [11]. In the model specified by Eq. (1) the atmospheric mixing angle at tree level is given by

$$\tan \theta_{\text{atm}} = \frac{\Lambda_2}{\Lambda_3} \quad (2)$$

$$\Lambda_i = \epsilon_i v_d + \mu v_i \quad (3)$$

where v_i are the sneutrino vevs and v_d is the vev of H_d^0 . It turns out that the dominant part of the $\tilde{\chi}_1^0$ -W- l_i coupling O_i^L is given by

$$O_i^L = \Lambda_i f(M_1, M_2, \mu, \tan \beta, v_d, v_u) \quad (4)$$

where the exact form of f can be found in Eq. (20) of ref. [11]. The important point is that f only depends on MSSM parameters but not on the R-parity violating parameters. Putting everything together one finds:

$$\tan^2 \theta_{\text{atm}} \simeq \left| \frac{\Lambda_2}{\Lambda_3} \right|^2 \simeq \frac{BR(\tilde{\chi}_1^0 \rightarrow \mu^\pm W^\mp)}{BR(\tilde{\chi}_1^0 \rightarrow \tau^\pm W^\mp)} \simeq \frac{BR(\tilde{\chi}_1^0 \rightarrow \mu^\pm \bar{q} q')}{BR(\tilde{\chi}_1^0 \rightarrow \tau^\pm \bar{q} q')}, \quad (5)$$

where the last equality is only approximate due to possible (small) contributions from three body decays of intermediate sleptons and squarks. The restriction to the hadronic final states of the W is necessary for the identification of the lepton flavour. Note that Eq. (5) is a prediction of the bilinear model independent of the R-parity conserving parameters.

2 Numerical Results

We take the SPS1a mSUGRA benchmark point [14] as a specific example, characterized by $m_0 = 100\text{GeV}$, $m_{\frac{1}{2}} = 250\text{GeV}$, $A_0 = -100\text{GeV}$, $\tan\beta = 10$, and $\text{sign}(\mu) = 1^1$. The low-energy parameters were derived using SPHENO 2.2 [15] and passed to PYTHIA 6.3 [16] using the recently defined SUSY Les Houches Accord [17]. The R-parity violating parameters (in MeV) at the low scale are given by: $\epsilon_1 = 43$, $\epsilon_2 = 100$, $\epsilon_3 = 10$, $v_1 = -2.9$, $v_2 = -6.7$ and $v_3 = -0.5$. For the neutrino sector we find $\Delta m_{\text{atm}}^2 = 3.8 \cdot 10^{-3} \text{ eV}^2$, $\tan^2 \theta_{\text{atm}} = 0.91$, $\Delta m_{\text{sol}}^2 = 2.9 \cdot 10^{-5} \text{ eV}^2$, $\tan^2 \theta_{\text{sol}} = 0.31$. Moreover, we find that the following neutralino branching ratios are larger than 1%:

$$\begin{aligned} \text{BR}(W^\pm \mu^\mp) &= 2.2\%, & \text{BR}(W^\pm \tau^\mp) &= 3.2\%, & \text{BR}(\bar{q}q' \mu^\mp) &= 1.5\%, \\ \text{BR}(\bar{q}q' \tau^\mp) &= 2.1\%, & \text{BR}(q\bar{q} \nu_i) &= 4.7\%, & \text{BR}(b\bar{b} \nu_i) &= 15.6\%, \\ \text{BR}(e^\pm \tau^\mp \nu_i) &= 5.9\%, & \text{BR}(\mu^\pm \tau^\mp \nu_i) &= 30.3\%, & \text{BR}(\tau^+ \tau^- \nu_i) &= 37.3\%, \end{aligned}$$

where we have summed over the neutrino final states as well as over the first two generations of quarks. Moreover, there are 0.2% of neutralinos decaying invisibly into three neutrinos. In the case that such events can be identified they can be used to distinguish this model from a model with trilinear R-parity violating couplings because in the latter case they are absent.

We now turn to the question to what extent the ratio, Eq. (5), could be measurable at an LHC experiment. The intention here is merely to illustrate the phenomenology and to give a rough idea of the possibilities. For simplicity, we employ a number of shortcuts; e.g. detector energy resolution effects are ignored and events are only generated at the parton level. Thus, we label a final-state quark or gluon which has $p_\perp > 15\text{GeV}$ and which lies within the fiducial volume of the calorimeter, $|\eta| < 4.9$, simply as ‘a jet’. Charged leptons are required to lie within the inner detector coverage, $|\eta| < 2.5$, and to have $p_\perp > 5\text{GeV}$ (electrons), $p_\perp > 6\text{GeV}$ (muons), or $p_\perp > 20\text{GeV}$ (taus). The assumed efficiencies for such leptons are [18] 75% for electrons, 95% for muons, and 85% for taus decaying in the 3-prong modes (we do not use the 1-prong decays), independent of p_\perp .

For SPS1a, the total SUSY cross section is $\sigma_{\text{SUSY}} \sim 41 \text{ pb}$. This consists mainly of gluino and squark pair production followed by subsequent cascades down to the LSP, the $\tilde{\chi}_1^0$. With an integrated luminosity of 100 fb^{-1} , approximately 8 million $\tilde{\chi}_1^0$ decays should thus have occurred in the detector.

An important feature of the scenario considered here is that the $\tilde{\chi}_1^0$ width is sufficiently small to result in a potentially observable displaced vertex. By comparing the decay length, $c\tau = 0.5 \text{ mm}$, with an estimated vertex resolution of about 20 microns in the transverse plane and 0.5 mm along the beam

¹Strictly speaking, the SPS points should be defined by their low-energy parameters as calculated with ISAJET 7.58.

mode	N_{gen}	ϵ_{rec}	$N_{rec}(100 \text{ fb}^{-1})$
$\tilde{\chi}_1^0 \rightarrow \mu W \rightarrow \mu q \bar{q}'$	235000	0.10	12500
$\tilde{\chi}_1^0 \rightarrow \tau W \rightarrow \tau_{3\text{-prong}} q \bar{q}'$	51600	0.054	1400

Table 1: Statistical sample, estimated reconstruction efficiencies, and expected event numbers.

axis, it is apparent that the two neutralino decay vertices should exhibit observable displacements in a fair fraction of events. Specifically, we require that both neutralino decays should occur outside an ellipsoid defined by 5 times the resolution. For at least one of the vertices (the ‘signal’ vertex), all three decay products ($\mu q \bar{q}'$ or $\tau q \bar{q}'$) must be reconstructed, while we only require one reconstructed decay product (jet in the inner detector or lepton in the inner detector whose track does not intersect the 5σ vertex resolution ellipsoid) for the second vertex (the ‘tag’ vertex).

Naturally, since the decay occurs within the detector, the standard SUSY missing E_{\perp} triggers are ineffective. Avoiding a discussion of detailed trigger menus (cf. [19]), we have approached the issue by requiring that each event contains either four jets, each with $p_{\perp} > 100\text{GeV}$, or two jets with $p_{\perp} > 100\text{GeV}$ together with a lepton (here meaning muon or electron) with $p_{\perp} > 20\text{GeV}$, or one jet with $p_{\perp} > 100\text{GeV}$ together with two leptons with $p_{\perp} > 20\text{GeV}$. Further, since the Standard Model background will presumably be dominated by $t\bar{t}$ events, we impose an additional parton-level b jet veto.

To estimate the efficiency with which decays into each channel can be reconstructed, a sample of 7.9 million SUSY events were generated with PYTHIA, and the above trigger and reconstruction cuts were imposed. To be conservative, we only include the resonant decay channels, where the quark pair at the signal vertex has the invariant mass of the W. The number of generated decays into each channel, the fractions remaining after cuts, and the expected total number of reconstructed events scaled to an integrated luminosity of 100 fb^{-1} are given in table 1. The comparatively small efficiencies owe mainly to the requirement that *both* neutralino decays should pass the 5σ vertex resolution cut. Nonetheless, using these numbers as a first estimate, the expected statistical accuracy of the ratio, $R = BR(\tilde{\chi}_1^0 \rightarrow \mu^{\pm} W^{\mp}) / BR(\tilde{\chi}_1^0 \rightarrow \tau^{\pm} W^{\mp})$, appearing in Eq. (5) becomes $\frac{\sigma(R)}{R} \simeq 0.028$.

3 Conclusions

We have studied neutralino decays in a model where bilinear R-parity violating terms are added to the usual MSSM Lagrangian. This model can successfully explain neutrino data and leads at the same time to *predictions*

for ratios of the LSP decay branching ratios. In particular we have considered a scenario where the lightest neutralino is the LSP. In this case the ratio $BR(\tilde{\chi}_1^0 \rightarrow \mu^\pm W^\mp)/BR(\tilde{\chi}_1^0 \rightarrow \tau^\pm W^\mp)$ is directly related to the atmospheric neutrino mixing angle. Provided R-parity violating SUSY is discovered, the measurement of this ratio at colliders would thus constitute an important test of the hypothesis of a supersymmetric origin of neutrino masses.

We have investigated the possibility of performing this measurement at a ‘generic’ LHC experiment, using PYTHIA to generate LHC SUSY events at the parton level and imposing semi-realistic acceptance and reconstruction cuts. Within this simplified framework, we find that the LHC should be sensitive to a possible connection between R-parity violating LSP decays and the atmospheric mixing angle, at least for scenarios with a fairly light sparticle spectrum and where the neutralino decay length is sufficiently large to give observable displaced vertices. Obviously, the numbers presented here represent crude estimates and should not be taken too literally. A more refined experimental analysis would be necessary for more definitive conclusions to be drawn.

Acknowledgements

We thank A. De Roeck and T. Sjöstrand for useful discussions. We are also grateful to the NorduGRID project for use of computing resources. W.P. is supported by the ‘Erwin Schrödinger fellowship No. J2272’ of the ‘Fonds zur Förderung der wissenschaftlichen Forschung’ of Austria and partly by the Swiss ‘Nationalfonds’.

References

- [1] Y. Fukuda et al. *Phys. Rev. Lett.*, 81:1562–1567, 1998.
- [2] S. Fukuda et al. *Phys. Rev. Lett.*, 86:5651–5655, 2001.
- [3] Q. R. Ahmad et al. *Phys. Rev. Lett.*, 87:071301, 2001.
- [4] K. Eguchi et al. *Phys. Rev. Lett.*, 90:021802, 2003.
- [5] Murray Gell-Mann, Pierre Ramond, and Richard Slansky. 1979. Print-80-0576 (CERN).
- [6] Tsutomu Yanagida. *Prog. Theor. Phys.*, 64:1103, 1980.
- [7] Rabindra N. Mohapatra and Goran Senjanovic. *Phys. Rev. Lett.*, 44:912, 1980.

-
- [8] J. C. Romão, M. A. Díaz, M. Hirsch, W. Porod, and J. W. F. Valle. *Phys. Rev.*, D61:071703, 2000.
 - [9] M. Hirsch, M. A. Díaz, W. Porod, J. C. Romão, and J. W. F. Valle. *Phys. Rev.*, D62:113008, 2000.
 - [10] M. A. Díaz, M. Hirsch, W. Porod, J. C. Romão, and J. W. F. Valle. *Phys. Rev.*, D68:013009, 2003.
 - [11] W. Porod, M. Hirsch, J. Romão, and J. W. F. Valle. *Phys. Rev.*, D63:115004, 2001.
 - [12] M. Hirsch, W. Porod, J. C. Romão, and J. W. F. Valle. *Phys. Rev.*, D66:095006, 2002.
 - [13] M. Hirsch and W. Porod. 2003. [hep-ph/0307364].
 - [14] B. C. Allanach et al. *Eur. Phys. J.*, C25:113–123, 2002.
 - [15] Werner Porod. *Comput. Phys. Commun.*, 153:275–315, 2003.
 - [16] Torbjorn Sjostrand, Leif Lonnblad, Stephen Mrenna, and Peter Skands. 2003. [hep-ph/0308153].
 - [17] P. Skands et al. 2003. [hep-ph/0311123].
 - [18] ATLAS Collaboration. CERN-LHCC-99-14.
 - [19] CMS Collaboration. CERN-LHCC-02-26.

Multiple Interactions and the Structure of Beam Remnants

Paper V

LU TP 04-08
hep-ph/0402078
March 2004

Multiple Interactions and the Structure of Beam Remnants

T. Sjöstrand¹ and P. Z. Skands²

*Department of Theoretical Physics,
Lund University,
Sölvegatan 14A,
S-223 62 Lund, Sweden*

Abstract

Recent experimental data have established some of the basic features of multiple interactions in hadron-hadron collisions. The emphasis is therefore now shifting, to one of exploring more detailed aspects. Starting from a brief review of the current situation, a next-generation model is developed, wherein a detailed account is given of correlated flavour, colour, longitudinal and transverse momentum distributions, encompassing both the partons initiating perturbative interactions and the partons left in the beam remnants. Some of the main features are illustrated for the Tevatron and the LHC.

¹torbjorn@thep.lu.se

²peter.skands@thep.lu.se

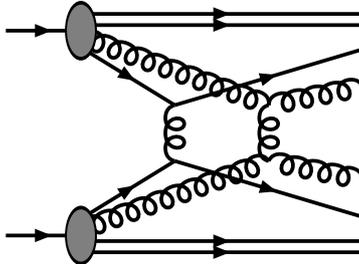


Figure 1: Schematic illustration of an event with two $2 \rightarrow 2$ perturbative interactions.

1 Introduction

The physics of high-energy hadron–hadron interactions has become a topic of increasing interest in recent years. With the Tevatron Run II well under way and with the startup of the LHC drawing closer, huge data samples are becoming available that will challenge our current understanding of this physics. From the point of view of QCD, many interesting questions remain to be answered, and we shall take up some of these in detail below. Moreover, for new physics searches and precision measurements, it is important that these questions can be given meaningful and trustworthy answers, since ever-present yet poorly-understood aspects of QCD can have a significant impact.

Much of the complexity involved in describing these phenomena — specifically the underlying event and minimum-bias collisions — derives from the composite nature of hadrons; we are dealing with objects which possess a rich internal structure that is not calculable from perturbation theory. This, however, does *not* imply that the physics of the underlying event as such has to be an inherently non-perturbative quagmire.

Viewing hadrons as ‘bunches’ of incoming partons, it is apparent that when two hadrons collide it is possible that several distinct pairs of partons collide with each other, as depicted in Fig. 1. Thus multiple interactions (also known as multiple scatterings) in hadronic collisions is a phenomenon which is a direct consequence of the composite nature of hadrons and which *must* exist, at some level. In fact, by extending simple perturbation theory to rather low p_{\perp} values, though still some distance above Λ_{QCD} , most inelastic events in high-energy hadronic collisions are guaranteed to contain several *perturbatively calculable* interactions [1]. Furthermore, such interactions — even when soft — can be highly important, causing non-trivial changes to the colour topology of the colliding system as a whole, with potentially drastic consequences for the particle multiplicity in the final state.

Nevertheless, traditionally the exploration of multiple interactions has not attracted much interest. For studies concentrating on high- p_{\perp} jets, perturbative QCD emission is a more important source of multijets than separate multiple interactions. The underlying event, on the other hand, has in this context often been viewed as a mess of soft QCD interactions, that cannot be described from first principles but is better simply parametrized.

However, such parametrizations, even while reasonably successful in describing the average underlying activity, are not sophisticated enough to adequately describe correlations and fluctuations. This relates for instance to jet profiles and jet pedestals, and to systematic as well as random shifts in jet energies. The lack of sophistication implies that, even when tuned to describe a few distributions well, they could not be trusted for extrapolations beyond the fit region. Hence, a sound understanding of multiple interactions is prerequisite for precision physics involving jets and/or the underlying event.

It is interesting to note that this can also impact physics studies in areas well beyond the conventional QCD ones. As an example, consider the search for a Higgs particle in the $h^0 \rightarrow \gamma\gamma$ channel at the LHC, where the Higgs mass resolution at high luminosity depends on picking the correct vertex between maybe 30 different pp events. If the event that contained the Higgs is special, by typically containing more charged particles (in general or in some regions of phase space), we would want to use that information to clean up the signal [2].

The crucial leap of imagination is to postulate that *all* particle production in inelastic hadronic collisions derives from the multiple-interactions mechanism. This is not to say that many nonperturbative and poorly known phenomena will not force their entrance on the stage, in going from the perturbative interactions to the observable hadrons, but the starting point is perturbative. If correct, this hypothesis implies that the typical Tevatron hadron-hadron collision may contain something like 2–6 interactions, and the LHC one 4–10.

A few models based on this picture were presented several years ago [1], and compared with the data then available. Though these models may still be tuned to give a reasonable description of the underlying event at various collider energies, several shortcuts had to be taken, particularly in the description of the nonperturbative aspects alluded to above. For instance, it was not possible to consider beam remnants with more than one valence quark kicked out.

The increased interest and the new data now prompts us to develop a more realistic framework for multiple interactions than the one in ref. [1], while making use of many of the same underlying ideas. This may not necessarily result in significant improvements for fits made to only a few distributions at a time, but we hope it will enhance our ability to simultaneously describe many different observations inside one framework, thereby improving the confidence with which we can make extrapolations from known measurements to new distributions and to higher energies.

One of the building blocks for the new model comes from our recent study of

baryon-number-violating processes [3]. We then had to address the hadronization of colour topologies of the same kind as found in baryons. Specifically, as an extension of the standard Lund string fragmentation framework [4], we explored the concept of a junction in which three string pieces meet, with a quark at the end of each string. This also opens the way to a more realistic description of multiple interaction events.

The resulting improvements, relative to the framework in [1], are especially notable in the description of the structure of the incoming hadrons, i.e. how flavours, colours, transverse and longitudinal momenta are correlated between all the partons involved, both those that undergo interactions and those that are left behind in the remnants. (Brief descriptions of some of these aspects can also be found in [5].) To give one specific example, we introduce parton densities that are modified according to the flavours already affected by interactions.

Clearly, the model we present here is not the final word. For instance, we defer for the future any discussions of processes involving photons in the initial state and whether and how diffractive topologies could arise naturally from several interactions with a net colour singlet exchange. More generally, the whole issue of colour correlations will require further studies. The model also allows some options in a few places. A reasonable range of possibilities can then be explored, and (eventually) experimental tests should teach us more about the course taken by nature.

This article is organized as follows. In Section 2 we give an introduction to multiple interactions in general, to the existing multiple interactions machinery, to other theoretical models, and to experimental data of relevance. Sections 3–5 then describe the improvements introduced by the current study: 3 a new option for impact-parameter dependence, 4 the main work on flavour and momentum space correlations, and 5 the very difficult topics of colour correlations and junction hadronization. Finally Section 6 provides some further examples of the resulting physics distributions and important tests, while Section 7 contains a summary and outlook.

2 Multiple Interactions Minireview

2.1 Basic Concepts

The cross section for QCD hard $2 \rightarrow 2$ processes, as a function of the p_{\perp}^2 scale, is given by

$$\frac{d\sigma_{\text{int}}}{dp_{\perp}^2} = \sum_{i,j,k} \int dx_1 \int dx_2 \int d\hat{t} f_i(x_1, Q^2) f_j(x_2, Q^2) \frac{d\hat{\sigma}_{ij \rightarrow kl}}{d\hat{t}} \delta\left(p_{\perp}^2 - \frac{\hat{t}\hat{u}}{\hat{s}}\right), \quad (1)$$

where $\hat{s} = x_1 x_2 s$. The jet cross section is twice as large, $\sigma_{\text{jet}} = 2\sigma_{\text{int}}$, since each interaction gives rise to two jets, to first approximation. In the following,

we will always refer to the interaction rather than the jet cross section, unless otherwise specified. We will also assume that the ‘hardness’ of processes is given by the p_{\perp} scale of the interaction, i.e. $Q^2 = p_{\perp}^2$.

The cross section for QCD $2 \rightarrow 2$ processes, the sum of $qq' \rightarrow qq'$, $q\bar{q} \rightarrow q'\bar{q}'$, $q\bar{q} \rightarrow gg$, $qg \rightarrow qg$, $gg \rightarrow gg$ and $gg \rightarrow q\bar{q}$, is dominated by t -channel gluon exchange contributions. (This justifies the use of ‘interaction’ and ‘scattering’ as almost synonymous.) In the $|\hat{t}| \ll \hat{s}$ limit, where $p_{\perp}^2 = \hat{t}\hat{u}/\hat{s} \approx |\hat{t}|$, quark and gluon interactions just differ by the colour factors, so approximately we may write

$$\frac{d\sigma_{\text{int}}}{dp_{\perp}^2} \approx \iint \frac{dx_1}{x_1} \frac{dx_2}{x_2} F(x_1, p_{\perp}^2) F(x_2, p_{\perp}^2) \frac{d\hat{\sigma}}{dp_{\perp}^2}, \quad (2)$$

where

$$\frac{d\hat{\sigma}}{dp_{\perp}^2} = \frac{8\pi\alpha_s^2(p_{\perp}^2)}{9p_{\perp}^4}, \quad (3)$$

and

$$F(x, Q^2) = \sum_q (xq(x, Q^2) + x\bar{q}(x, Q^2)) + \frac{9}{4}xg(x, Q^2). \quad (4)$$

Thus, for constant α_s and neglecting the x integrals, the integrated cross section above some $p_{\perp\text{min}}$ is divergent in the limit $p_{\perp} \rightarrow 0$:

$$\sigma_{\text{int}}(p_{\perp\text{min}}) = \int_{p_{\perp\text{min}}}^{\sqrt{\hat{s}}/2} \frac{d\sigma}{dp_{\perp}} dp_{\perp} \propto \frac{1}{p_{\perp\text{min}}^2}. \quad (5)$$

A numerical illustration of this divergence is given in Fig. 2. Note that the actual fall-off is everywhere steeper than $1/p_{\perp}^2$. We have here used full $2 \rightarrow 2$ matrix elements and the CTEQ 5L parton density parametrizations [6], which are valid for $Q > 1.1$ GeV and $x > 10^{-6}$; therefore results at the lowest p_{\perp} values are not to be taken too literally. For the studies in this article we are basing ourselves on leading-order cross sections and parton densities, with nontrivial higher-order corrections only approximately taken into account by the addition of parton showers. Nevertheless, the trend is quite clear, with an integrated cross section that exceeds the total $pp/\bar{p}p$ cross section σ_{tot} (in the parametrization of ref. [7]) for $p_{\perp\text{min}}$ of the order of a few GeV. As already mentioned in the introduction, this is well above Λ_{QCD} , so one cannot postulate a breakdown of perturbation theory in the conventional sense.

The resolution of the $\sigma_{\text{int}} > \sigma_{\text{tot}}$ paradox probably comes in two steps.

Firstly, the interaction cross section is an inclusive number. Thus, if an event contains two interactions it counts twice in σ_{int} but only once in σ_{tot} , and so on for higher multiplicities. Thereby we may identify $\langle n \rangle(p_{\perp\text{min}}) = \sigma_{\text{int}}(p_{\perp\text{min}})/\sigma_{\text{tot}}$ with the average number of interactions above $p_{\perp\text{min}}$ per event, and that number may well be above unity.

One of the problems we will consider further in this article is that this simple calculation of $\langle n \rangle(p_{\perp\text{min}})$ does not take into account energy-momentum conservation effects. Specifically, the average \hat{s} of a scattering decreases slower with

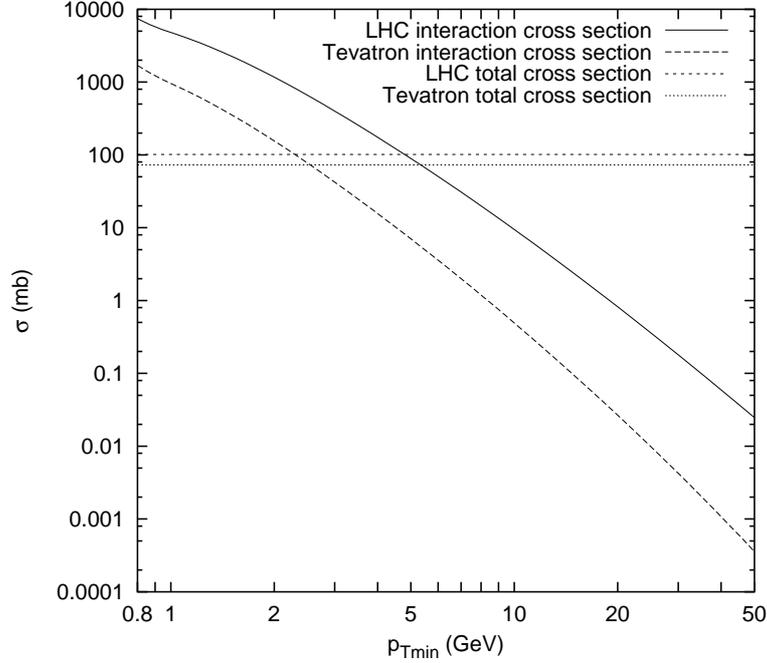


Figure 2: The integrated interaction cross section σ_{int} above $p_{\perp\text{min}}$ for the Tevatron, with 1.8 TeV $p\bar{p}$ collisions, and the LHC, with 14 TeV pp ones. For comparison, the flat lines represent the respective total cross section.

$p_{\perp\text{min}}$ than the number of interactions increases, so naively the total amount of scattered partonic energy becomes infinite. Thus corrections reduce the $\langle n \rangle(p_{\perp\text{min}})$ number, but not sufficiently strongly: one is led to a picture with too little of the incoming energy remaining in the small-angle beam jet region [1].

Secondly, a more credible reason for taming the rise of $\langle n \rangle(p_{\perp\text{min}})$ is that the incoming hadrons are colour singlet objects. Therefore, when the p_{\perp} of an exchanged gluon is made small and the transverse wavelength correspondingly large, the gluon can no longer resolve the individual colour charges, and the effective coupling is decreased. Note that perturbative QCD calculations are always performed assuming free incoming and outgoing quark and gluon states, rather than partons inside hadrons, and thus do not address this kind of nonperturbative screening effects.

A naive estimate of an effective lower cutoff would be

$$p_{\perp\text{min}} \simeq \frac{\hbar}{r_p} \approx \frac{0.2 \text{ GeV} \cdot \text{fm}}{0.7 \text{ fm}} \approx 0.3 \text{ GeV} \simeq \Lambda_{\text{QCD}} , \quad (6)$$

but this again appears too low. The proton radius r_p has to be replaced by the typical colour screening distance d , i.e. the average size of the region within which the net compensation of a given colour charge occurs. This number is not known from first principles, so effectively one is forced to introduce some kind of cutoff parameter, which can then just as well be put in transverse momentum space. The simplest choice is to introduce a step function $\theta(p_\perp - p_{\perp\min})$, such that the perturbative cross section completely vanishes below the $p_{\perp\min}$ scale. A more realistic alternative is to note that the jet cross section is divergent like $\alpha_s^2(p_\perp^2)/p_\perp^4$, and that therefore a factor

$$\frac{\alpha_s^2(p_{\perp 0}^2 + p_\perp^2)}{\alpha_s^2(p_\perp^2)} \frac{p_\perp^4}{(p_{\perp 0}^2 + p_\perp^2)^2} \quad (7)$$

would smoothly regularize the divergences, now with $p_{\perp 0}$ as the free parameter to be tuned to data. Empirically the two procedures give similar numbers, $p_{\perp\min} \approx p_{\perp 0}$, and both of the order of 2 GeV.

The parameters do not have to be energy-independent, however. Higher energies imply that parton densities can be probed at smaller x values, where the number of partons rapidly increases. Partons then become closer packed and the colour screening distance d decreases. Just like the small- x rise goes like some power of x one could therefore expect the energy dependence of $p_{\perp\min}$ and $p_{\perp 0}$ to go like some power of CM energy. Explicit toy simulations [8] lend some credence to such an ansatz, although with large uncertainties. Alternatively, one could let the cutoff increase with decreasing x ; this would lead to a similar phenomenology since larger energies probe smaller x values.

2.2 Our Existing Models

The models developed in ref. [1] have been implemented and are available in the PYTHIA event generator. They form the starting point for the refinements we will discuss further on, so we here review some of the main features.

The approach is not intended to cover elastic or diffractive physics, so the $\sigma_{\text{int}}(p_{\perp\min}, s)$ or $\sigma_{\text{int}}(p_{\perp 0}, s)$ interaction cross section is distributed among the $\sigma_{\text{nd}}(s)$ nondiffractive inelastic one [7, 10]. The average number of interactions per such event is then the ratio $\langle n \rangle = \sigma_{\text{int}}/\sigma_{\text{nd}}$. As a starting point we will assume that all hadron collisions are equivalent, i.e. without an impact parameter dependence, and that the different parton-parton interactions take place independently of each other. The number of interactions per event is then distributed according to a Poisson distribution with mean $\langle n \rangle$, $\mathcal{P}_n = \langle n \rangle^n \exp(-\langle n \rangle)/n!$.

One (not used) approach would be, for each new event, to pick the actual number of interactions n according to the Poissonian, and select the n p_\perp values independently according to eq. (1). One disadvantage is that this does not take into account correlations, even such basic ones as energy-momentum

conservation: the sum of interaction energies may well exceed the total CM energy.

In an event with several interactions, it is therefore convenient to impose an ordering. The logical choice is to arrange the scatterings in falling sequence of p_{\perp} values. The ‘first’ scattering is thus the hardest one, with the ‘subsequent’ (‘second’, ‘third’, etc.) successively softer. This terminology is not primarily related to any picture in physical time although, by the uncertainty relation, large momentum transfers implies short timescales. When averaging over all configurations of soft partons, one should effectively obtain the standard QCD phenomenology for a hard scattering, e.g. in terms of parton distributions. Correlation effects, known or estimated, can be introduced in the choice of subsequent scatterings, given that the ‘preceding’ (harder) ones are already known. This will be developed further in Section 4.

The generation of a sequence $\sqrt{s}/2 > p_{\perp 1} > p_{\perp 2} > \dots > p_{\perp n} > p_{\perp \min}$ now becomes one of determining $p_{\perp i}$ from a known $p_{\perp i-1}$, according to the probability distribution

$$\frac{d\mathcal{P}}{dp_{\perp i}} = \frac{1}{\sigma_{\text{nd}}} \frac{d\sigma}{dp_{\perp}} \exp \left[- \int_{p_{\perp}}^{p_{\perp i-1}} \frac{1}{\sigma_{\text{nd}}} \frac{d\sigma}{dp'_{\perp}} dp'_{\perp} \right]. \quad (8)$$

The exponential expression is the ‘form factor’ from the requirement that no interactions occur between $p_{\perp i-1}$ and $p_{\perp i}$, cf. radioactive decays or the Sudakov form factor [11] of parton showers.

When used with the standard differential cross section $d\sigma/dp_{\perp}$, eq. (8) gives the same Poisson distribution as above. This time n is not known beforehand, but is defined by the termination of the iterative procedure. Now, however, $d\sigma/dp_{\perp}$ can be modified to take into account the effects of the $i-1$ preceding interactions. Specifically, parton distributions are not evaluated at x_i for the i ’th scattered parton from a hadron, but at the rescaled value

$$x'_i = \frac{x_i}{1 - \sum_{j=1}^{i-1} x_j}, \quad (9)$$

so that it becomes impossible to scatter more energy than initially available in the incoming beam. This also dynamically suppresses the high-multiplicity tail of the Poissonian and thereby reduces the average number of interactions.

In a fraction of the events studied, there will be no hard scattering at all above $p_{\perp \min}$. Such events are associated with nonperturbative low- p_{\perp} physics, and are simulated by exchanging a very soft gluon between the two colliding hadrons, making the hadron remnants colour-octet objects rather than colour-singlet ones. If only valence quarks are considered, the colour-octet state of a baryon can be decomposed into a colour triplet quark and an antitriplet diquark. In a baryon–baryon collision, one would then obtain a two-string picture, with each string stretched from the quark of one baryon to the diquark

of the other. A baryon–antibaryon collision would give one string between a quark and an antiquark and another one between a diquark and an antidiquark.

In a hard interaction, the number of possible string drawings are many more, and the overall situation can become quite complex. In the studies preceding this work, several simplifications were made. The hardest interaction was selected with full freedom of flavour choice and colour topology, but for the subsequent ones only three simple recipes were available:

- Interactions of the $gg \rightarrow gg$ type, with the two gluons in a colour-singlet state, such that a double string is stretched directly between the two outgoing gluons, decoupled from the rest of the system.
- Interactions $gg \rightarrow gg$, but colour correlations assumed to be such that each of the gluons is connected to one of the strings ‘already’ present. Among the different possibilities of connecting the colours of the gluons, the one which minimizes the total increase in string length is chosen. This is in contrast to the previous alternative, which roughly corresponds to a maximization (within reason) of the extra string length.
- Interactions $gg \rightarrow q\bar{q}$, with the final pair again in a colour-singlet state, such that a single string is stretched between the outgoing q and \bar{q} .

The three possibilities can be combined in suitable fractions.

Many further approximations were also required in the old framework, e.g. the addition of initial- and final-state parton showers was feasible only for the hardest interaction, and we address several of those in the following.

The model also includes several options for the impact-parameter dependence. This offers an additional element of variability: central collisions on the average will contain more interactions than peripheral ones. Even if a Poisson distribution in the number of interactions would be assumed for each impact parameter separately, the net result would be a broader-than-Poisson distribution. The amount of broadening can be ‘tuned’ by the choice of impact-parameter profile. We discuss this further in section 3, where a new set of profiles is studied.

The above framework was originally formulated for $p\bar{p}/pp$ collisions, but has also been extended to γp and $\gamma\gamma$ interactions [12]. In these latter processes, however, the nature of the photon needs to be modelled in detail, and this introduces many further uncertainties. A study of such aspects is beyond the scope of the current article.

2.3 Other Models

While the models of ref. [1] may well be the ones most frequently studied, owing to their implementation in PYTHIA [9], a number of other models also exist.

Many of the basic concepts have also been studied separately. We here give a few examples, without any claim of completeness.

In Dual Topological Unitarization (DTU) language [13], and the Dual Parton Model based on it [14], or other similar techniques [15], inelastic events are understood in terms of cut pomerons [16]. Translated into modern terminology, each cut pomeron corresponds to the exchange of a soft gluon, which results in two ‘strings’ being drawn between the two beam remnants. Uncut pomerons give virtual corrections that help preserve unitarity. A variable number of cut pomerons are allowed. This approach has been the basis for the simulation of underlying events in ISAJET [17], and was the starting point for DTUJET [18]. However, note that cut pomerons were originally viewed as purely soft objects, and so did not generate any transverse momentum, unlike the multiple interactions considered in this article. In DTUJET and its PHOJET [19] and DP-MJET [20] relatives, however, also hard interactions have been included, so that the picture now is one of both hard and soft pomerons, ideally with a smooth transition between the two. Since the three related programs make use of the PYTHIA hadronization description, the differences relative to our scenarios is more a matter of details (but “the devil is in the details”) than of any basic incompatibility.

The possibility of observing two separate hard interactions has been proposed since long [21], and from that has also developed a line of studies on the physics framework for having several hard interactions [22], also involving e.g. electroweak processes [23]. Again this is similar to what we do, except that lower p_{\perp} values and the transition to nonperturbative physics are not normally emphasized.

The possibility of multiple interactions has also been implicit [24] or explicit [25] in many calculations of total cross sections for hadron–hadron, hadron– γ and $\gamma\gamma$ events. The increase of σ_{int} with CM energy is here directly driving an increase also of σ_{tot} ; that the latter is rising slower than the former comes out of an eikonalization procedure that implies also an increasing $\langle n_{\text{int}} \rangle$.

Multiple interactions require an ansatz for the structure of the incoming beams, i.e. correlations between the constituent partons. Some of these issues have been studied, e.g. with respect to longitudinal momenta [14, 26], colours [27] or impact parameter [28], but very little of this has been tested experimentally. Dense-packing of partons could become an issue [15], of unknown importance, but not necessarily a major one [29].

The HERWIG [30] event generator does not contain any physics simulation of multiple interactions. Instead a parametrization procedure originally suggested by UA5 [31] is used, without any underlying physics scenario. It thus does parametrize multiplicity and rapidity distributions, but does not contain any minijet activity in the underlying event. The add-on JIMMY package [32] offers a multiple-interaction component (both for pp, γp and $\gamma\gamma$ events), which has recently been extended to include also a model of soft interactions [33].

The introduction of unintegrated parton densities, as used in the BFKL/CCFM/LDC approaches to initial-state radiation [34–36], allows the possibility to replace our $p_{\perp\text{min}}/p_{\perp 0}$ cutoff by parton densities that explicitly vanish in the $p_{\perp} \rightarrow 0$ limit [37]. This opens up the possibility of an alternative implementation of multiple interactions [38].

In heavy-ion collisions the multiple interactions rate can become huge [39]. For small impact parameters, a major fraction of the energy of the two incoming nuclei is carried by partons undergoing perturbative interactions, and which therefore define a ‘resolved’ partonic content. This suggests a mechanism for the construction of an ‘initial state’ for the continued formation and thermalization (or not) of a quark–gluon plasma.

2.4 Experimental Tests

Experimental input to the understanding of multiple interactions comes in essentially three categories: direct observation of double parton scattering, event properties that directly and strongly correlate with multiple interactions, and event properties that do not point to multiple interactions by themselves but still constrain multiple interactions models.

If an event contains two uncorrelated $2 \rightarrow 2$ interactions, we expect to find four jets, grouped in two pairs that each internally have roughly opposite and compensating transverse momenta, and where the relative azimuthal angle between the scattering planes is isotropic. Neither of these properties are expected in a $2 \rightarrow 4$ event, where two of the partons can be thought of as bremsstrahlung off a basic $2 \rightarrow 2$ process. The problem is that $2 \rightarrow 4$ processes win out at large p_{\perp} , so there is a delicate balance between having large enough jet p_{\perp} that the jets can be well measured and still not so large that the signal drowns.

When the $p_{\perp\text{min}}$ of the jets is sufficiently large that $\exp(-\langle n \rangle) \approx 1$, Poisson statistics implies that $\mathcal{P}_2 = \mathcal{P}_1^2/2$, where \mathcal{P}_i is the probability to have i interactions. Traditionally this is rewritten as

$$\sigma_2 = \sigma_{\text{nd}} \left(\frac{\sigma_1}{\sigma_{\text{nd}}} \right)^2 \frac{1}{2} \frac{\sigma_{\text{nd}}}{\sigma_{\text{eff}}} = \frac{1}{2} \frac{\sigma_1^2}{\sigma_{\text{eff}}}, \quad (10)$$

where the ratio $\sigma_{\text{nd}}/\sigma_{\text{eff}}$ gauges deviations from the Poissonian ansatz. Values above unity, i.e. $\sigma_{\text{eff}} < \sigma_{\text{nd}}$, arise naturally in models with variable impact parameter.

The first observation of double parton scattering is by AFS [40]. The subsequent UA2 study [41] ends up quoting an upper limit, but has a best fit that requires them. A CDF study [42] also found them. These experiments all had to contend with limited statistics and uncertain background estimates. The strongest signal has been obtained with a CDF study involving three jets and a hard photon [43]; here $\sigma_2 = \sigma_A \sigma_B / \sigma_{\text{eff}}$, without a factor 1/2, since the two

$2 \rightarrow 2$ processes A and B are inequivalent. In all cases, including the UA2 best fit, σ_{eff} comes out smaller than σ_{nd} ; typically the double parton scattering cross section is a factor of three to four larger than the Poissonian prediction. For instance, the CDF number is $\sigma_{\text{eff}} = 14.5 \pm 1.7^{+1.7}_{-2.3}$ mb. More recently, ZEUS has observed a signal in γp events [44]. The D0 four-jet study shows the need to include multiple interactions, but this is not quantified [45].

So far, no direct tests of triple or more parton scattering exist. However, the UA1 minijet rates [46], going up to 5 jets, are difficult to understand without such events.

Tests involving jets at reasonably large p_{\perp} values do not probe the total rate of multiple interactions, since the bulk of the interactions occur at so small p_{\perp} values that they cannot be identified as separate jets. By the way colours are drawn across the event, soft partons can drive the production of particles quite out of proportion to the p_{\perp} values involved, however. The multiplicity distribution of multiple interactions thereby strongly influences the multiplicity distribution of charged hadrons, n_{ch} .

A notable aspect here is that the measured n_{ch} distribution, when expressed in the KNO variable $z = n_{\text{ch}}/\langle n_{\text{ch}} \rangle$ [47], is getting broader with increasing CM energy [48, 49]. This is contrary to the essentially Poissonian hadronization mechanism of the string model, where the KNO distribution becomes narrower. As an example, consider the UA5 measurements at 900 GeV [48], where $\langle n_{\text{ch}} \rangle = 35.6$ and $\sigma(n_{\text{ch}}) = 19.5$, while the Poissonian prediction would be $\sigma(n_{\text{ch}}) = \sqrt{\langle n_{\text{ch}} \rangle} = 6.0$. It is possible to derive approximate KNO scaling in e^+e^- annihilation [50], but this then rests on having a perturbative shower that involves the whole CM energy. However, allowing for at most one interaction in $p\bar{p}$ events and assuming that hadronization is universal (so it can be tuned to e^+e^- data), there is no (known) way to accommodate the experimental multiplicity distributions, neither the rapid increase with energy of the average nor the large width. Either hadronization is *very* different in hadronic events from e^+e^- ones, or one must accept multiple interactions as a reality. (Unfortunately it is difficult to test the ‘hadronization universality’ hypothesis completely separated from the multiple interactions and other assumptions. To give two examples, the relative particle flavour composition appears to be almost but not quite universal [51], and low-mass diffractive events display ‘string-like’ flavour correlations [52].)

Further support is provided by the study of forward–backward multiplicity correlations. For instance, UA5 and E735 define a ‘forward’ n_F and a ‘backward’ n_B multiplicity in pseudorapidity windows of one unit each, separated by a $\Delta\eta$ variable-width gap in the middle [53]. A forward–backward correlation strength is now defined by

$$b = \frac{\langle (n_F - \langle n_F \rangle) (n_B - \langle n_B \rangle) \rangle}{\sigma(n_F) \sigma(n_B)} = \frac{\langle n_F n_B \rangle - \langle n_F \rangle^2}{\langle n_F^2 \rangle - \langle n_F \rangle^2}, \quad (11)$$

where the last equality holds for a symmetric η distribution, i.e. for pp/p \bar{p} but not for γ p. Measurements give a positive and surprisingly large b , also for $\Delta\eta$ of several units of rapidity. So it appears that there is some global quantity, different for each event, that strongly influences particle production in the full phase space. Again known fragmentation mechanisms are too local, and effects of a single hard interaction not strong enough. But the number of multiple interactions is indeed a global quantity of the desired kind, and multiple-interaction models can describe the data quite well.

It is a matter of taste which evidence is valued highest. The direct observation of double parton scattering is easily recognized as evidence for the multiple-interactions concept, but only affects a tiny fraction of the cross section. By comparison, the broadening multiplicity distribution and the strong forward-backward correlations offer more indirect evidence, but ones strongly suggesting that the bulk of events have several interactions. We are not aware of any realistic alternative explanations for either of the observables.

Another interesting phenomenon is the pedestal effect: events with high- p_{\perp} jets on the average contain more underlying activity than minimum-bias ones, also well away from jets themselves. It has been observed by several collaborations, like UA1 [54], CDF [55,56] and H1 [57]. When the jet energy is varied from next-to-nothing to the highest possible, the underlying activity initially increases, but then flattens out around $p_{\perp\text{jet}} = 10$ GeV (details depending on the jet algorithm used and the CM energy). This fits very nicely with an impact-parameter-dependent multiple-interactions scenario: the presence of a higher p_{\perp} scale biases events towards a smaller impact parameter and thereby a higher additional activity, but once $\sigma_{\text{int}}(p_{\perp\text{jet}}) \ll \sigma_{\text{nd}}$ the bias saturates [1]. The height of the pedestal depends on the form of the overlap function $\mathcal{O}(b)$, and can thus be adjusted, while the $p_{\perp\text{jet}}$ at which saturation occurs is rather model-insensitive, and in good agreement with the data.

The presence of pedestals also affects all measurements of jet profiles [55,59]. It can lead to seemingly broader jets, when the full underlying event cannot be subtracted, and enrich the jet substructure, when a multiple-interactions jet is mistaken for radiation off the hard subprocess. It can also affect (anti)correlations inside a jet and with respect to the rest of the event [57].

Many further observables influence the modeling and understanding of multiple interactions, without having an immediate interpretation in those terms. A notable example here is the $\langle p_{\perp} \rangle(n_{\text{ch}})$ distribution, i.e. how the average transverse momentum of charged particles varies as a function of the total charged multiplicity. The observed increasing trend [58] is consistent with multiple interactions: large multiplicity implies many interactions and therefore more perturbatively generated p_{\perp} to be shared between the hadrons. For it to work, however, each new interaction should add proportionately less to the total n_{ch} than to the total p_{\perp} . Whether this is the case strongly depends on the colour connections between the interactions, i.e. whether strings tend to con-

nect nearest neighbours in momentum space or run criss-cross in the event. A rising trend can easily be obtained, but it is a major challenge to get the quantitative behaviour right, as we shall see.

Finally, one should mention that global fits to hadron collider data [60–64] clearly point to the importance of a correct understanding of multiple interactions, and constrains models down to rather fine details. This brings together many of the aspects raised above, plus some further ones. A convenient reference for our continued discussion is Tune A, produced by R.D. Field, which is known to describe a large set of CDF minimum bias and jet data [61]. Relative to the defaults of the old scenario, it assumes $p_{\perp 0} = 2.0$ GeV (PARP(82)=2.0) at the reference energy 1.8 TeV (PARP(89)=1800.0), with an energy rescaling proportional to $E_{\text{cm}}^{1/4}$ (PARP(90)=0.25). It is based on a double Gaussian matter distribution (MSTP(82)=4), with half of the matter (PARP(83)=0.5) in a central core of radius 40% of the rest (PARP(84)=0.4). Almost all of the subsequent interactions are assumed to be of the type $gg \rightarrow gg$ with minimal string length (PARP(85)=0.9, PARP(86)=0.95). Finally the matching of the initial-state showers to the hard scattering is done at a scale $Q_{\text{shower}}^2 = 4p_{\perp \text{hard}}^2$ (PARP(67)=4.0).

The above parameter set is sensible, within the framework of the model, although by no means obvious. The matter distribution is intermediate between the extremes already considered in [1], while the string drawing is more biased towards small string lengths than foreseen there. The $p_{\perp 0}$ energy dependence is steeper than previously used, but in a sensible range, as follows. In reggeon theory, a Pomeron intercept of $1 + \epsilon$ implies a total cross section rising like s^ϵ , and a small- x gluon density like $xg(x) \propto x^{-\epsilon}$ (at small Q^2). A $p_{\perp 0}$ rising (at most) like s^ϵ would then be acceptable, while one rising significantly faster would imply a decreasing interaction cross section $\sigma_{\text{int}}(p_{\perp 0})$, and by implication a decreasing σ_{tot} , in contradiction with data. The DL fit to σ_{tot} [7] gives $\epsilon \approx 0.08$, which would imply (at most) a $p_{\perp 0}$ dependence like $s^{0.08} = E_{\text{cm}}^{0.16}$. However, σ_{tot} already represents the unitarization of multiple-pomeron exchanges, and the ‘bare’ pomeron intercept should be larger than this, exactly by how much being a matter of some debate [65]. A value like $\epsilon_{\text{bare}} \approx 0.12$ is here near the lower end of the sensible range; the $xg(x)$ shape is consistent with a rather larger value. Since it is actually the bare pomeron that corresponds to a single interaction, an $E_{\text{cm}}^{0.25}$ behaviour is thereby acceptable.

3 Impact-Parameter Dependence

In the simplest multiple-interactions scenarios, it is assumed that the initial state is the same for all hadron collisions. More realistically, one should include the possibility that each collision also could be characterized by a varying impact parameter b [1]. Within the classical framework we use here, b is to be

thought of as a distance of closest approach, not as the Fourier transform of the momentum transfer. A small b value corresponds to a large overlap between the two colliding hadrons, and hence an enhanced probability for multiple interactions. A large b , on the other hand, corresponds to a grazing collision, with a large probability that no parton-parton interactions at all take place.

In order to quantify the concept of hadronic matter overlap, one may assume a spherically symmetric distribution of matter inside a hadron at rest, $\rho(\mathbf{x}) d^3x = \rho(r) d^3x$. For simplicity, the same spatial distribution is taken to apply for all parton species and momenta. Several different matter distributions have been tried. A Gaussian ansatz makes the subsequent calculations especially transparent, but there is no reason why this should be the correct form. Indeed, it appears to lead to a somewhat too narrow multiplicity distribution and too little of a pedestal effect. The next simplest choice, that does provide more fluctuations, is a double Gaussian

$$\rho(r) \propto \frac{1-\beta}{a_1^3} \exp\left\{-\frac{r^2}{a_1^2}\right\} + \frac{\beta}{a_2^3} \exp\left\{-\frac{r^2}{a_2^2}\right\}. \quad (12)$$

This corresponds to a distribution with a small core region, of radius a_2 and containing a fraction β of the total hadronic matter, embedded in a larger hadron of radius a_1 . If we want to give a deeper meaning to this ansatz, beyond it containing two more adjustable parameters, we could imagine it as an intermediate step towards a hadron with three disjoint core regions ('hot spots'), reflecting the presence of three valence quarks, together carrying the fraction β of the proton momentum. One could alternatively imagine a hard hadronic core surrounded by a pion cloud. Such details would affect e.g. the predictions for the t distribution in elastic scattering, but are not of any consequence for the current topics.

For a collision with impact parameter b , the time-integrated overlap $\mathcal{O}(b)$ between the matter distributions of the colliding hadrons is given by

$$\begin{aligned} \mathcal{O}(b) &\propto \int dt \int d^3x \rho(x, y, z) \rho(x+b, y, z+t) \\ &\propto \frac{(1-\beta)^2}{2a_1^2} \exp\left\{-\frac{b^2}{2a_1^2}\right\} \\ &\quad + \frac{2\beta(1-\beta)}{a_1^2 + a_2^2} \exp\left\{-\frac{b^2}{a_1^2 + a_2^2}\right\} \\ &\quad + \frac{\beta^2}{2a_2^2} \exp\left\{-\frac{b^2}{2a_2^2}\right\}. \end{aligned} \quad (13)$$

The necessity to use boosted $\rho(\mathbf{x})$ distributions has been circumvented by a suitable scale transformation of the z and t coordinates. The overlap function $\mathcal{O}(b)$ is closely related to the $\Omega(b)$ of eikonal models [24], but is somewhat simpler in spirit.

The larger the overlap $\mathcal{O}(b)$ is, the more likely it is to have interactions between partons in the two colliding hadrons. In fact, to first approximation, there should be a linear relationship

$$\langle \tilde{n}(b) \rangle = k\mathcal{O}(b) , \quad (14)$$

where $\tilde{n} = 0, 1, 2, \dots$ counts the number of interactions when two hadrons pass each other with an impact parameter b . At this stage k is an undefined constant of proportionality, to be specified below.

For each given impact parameter, the number of interactions \tilde{n} is assumed to be distributed according to a Poissonian, $\mathcal{P}_{\tilde{n}} = \langle \tilde{n} \rangle^{\tilde{n}} \exp(-\langle \tilde{n} \rangle) / \tilde{n}!$, before energy-momentum and other constraints are included. If the matter distribution has a tail to infinity (as the Gaussians do), events may be obtained with arbitrarily large b values. In order to obtain finite total cross sections, it is necessary to assume that each event contains at least one semi-hard interaction. (Unlike the simpler, impact-parameter-independent approach above, where $p_{\perp} = 0$ no-interaction events are allowed as a separate class.) The probability that two hadrons, passing each other with an impact parameter b , will produce a real event is then given by

$$\mathcal{P}_{\text{int}}(b) = \sum_{\tilde{n}=1}^{\infty} \mathcal{P}_{\tilde{n}}(b) = 1 - \mathcal{P}_0(b) = 1 - \exp(-\langle \tilde{n}(b) \rangle) = 1 - \exp(-k\mathcal{O}(b)) , \quad (15)$$

according to Poisson statistics. The average number of interactions per event at impact parameter b is now $\langle n(b) \rangle = \langle \tilde{n}(b) \rangle / \mathcal{P}_{\text{int}}(b)$, where the denominator comes from the removal of hadron pairs that pass without interaction, i.e. which do not produce any events. While the removal of $\tilde{n} = 0$ from the potential event sample gives a narrower-than-Poisson interaction distribution at each fixed b , the variation of $\langle n(b) \rangle$ with b gives a b -integrated broader-than-Poisson interaction multiplicity distribution.

Averaged over all b the relationship $\langle n \rangle = \sigma_{\text{int}} / \sigma_{\text{nd}}$ should still hold. Here, as before, σ_{int} is the integrated interaction cross section for a given regularization prescription at small p_{\perp} , while the inelastic nondiffractive cross section σ_{nd} is taken from parametrizations [7, 10]. This relation can be used to solve for the proportionality factor k in eq. (14). Note that, since now each event has to have at least one interaction, $\langle n \rangle > 1$, one must ensure that $\sigma_{\text{int}} > \sigma_{\text{nd}}$. The $p_{\perp 0}$ parameter has to be chosen accordingly small — since now the concept of no-interaction low- p_{\perp} events is gone, aesthetically it is more appealing to use the smooth $p_{\perp 0}$ turnoff than the sharp $p_{\perp \text{min}}$ cutoff, and thereby populate interactions continuously all the way down to $p_{\perp} = 0$. The whole approach can be questioned at low energies, since then very small $p_{\perp 0}$ values would be required, so that many of the interactions would end up in the truly nonperturbative p_{\perp} region.

Technically, the combined selection of b and a set of scattering $p_{\perp i}$ values now becomes more complicated [1, 9]. It can be reduced to a combined choice

of b and $p_{\perp 1}$, according to a generalization of eq. (8)

$$\frac{d\mathcal{P}}{dp_{\perp 1} d^2b} = \frac{\mathcal{O}(b)}{\langle \mathcal{O} \rangle} \frac{1}{\sigma_{\text{nd}}} \frac{d\sigma}{dp_{\perp}} \exp \left[-\frac{\mathcal{O}(b)}{\langle \mathcal{O} \rangle} \int_{p_{\perp}}^{\sqrt{s}/2} \frac{1}{\sigma_{\text{nd}}} \frac{d\sigma}{dp'_{\perp}} dp'_{\perp} \right]. \quad (16)$$

The removal of the $\tilde{n} = 0$ non-events leads to a somewhat special definition of the average [1]

$$\langle \mathcal{O} \rangle = \frac{\int \mathcal{O}(b) d^2b}{\int \mathcal{P}_{\text{int}}(b) d^2b} = \frac{1}{k} \frac{\sigma_{\text{int}}}{\sigma_{\text{nd}}}. \quad (17)$$

The subsequent interactions can be generated sequentially in falling p_{\perp} as before, with the only difference that $d\sigma/dp_{\perp}^2$ now is multiplied by $\mathcal{O}(b)/\langle \mathcal{O} \rangle$, where b is fixed at the value chosen above.

Note that this lengthy procedure, via $\rho(r)$ and $\mathcal{O}(b)$, is not strictly necessary: the probability \mathcal{P}_n for having n interactions could be chosen according to any desired distribution. However, with only \mathcal{P}_n known and an n selected from this distribution, there is no obvious way to order the interactions in p_{\perp} during the generation stage, with lower- p_{\perp} interactions modified by the flavours, energies and momenta of higher- p_{\perp} ones. (This problem is partly addressed in ref. [38], by a post-facto ordering of interactions and a subsequent rejection of some of the generated interactions, but flavour issues are not easily solved that way.)

There is also another issue, the parton-level pedestal effect, related to the transition from hard events to soft ones. To first approximation, the likelihood that an event contains a very hard interaction is proportional to $n\mathcal{P}_n$, since n interactions in an event means n chances for a hard one. If the requested hardest p_{\perp} is gradually reduced, the bias towards large n dies away and turns into its opposite: for events with the hardest $p_{\perp} \rightarrow 0$ the likelihood of further interactions vanishes. The interpolation between these two extremes can be covered if an impact parameter is chosen, and thereby an $\mathcal{O}(b)$, such that one can calculate the probability of *not* having an interaction harder than the requested hardest one, i.e. the exponential in eq. (16).

If the Gaussian matter distribution is the simplest possible choice, the double Gaussian in some respects is the next-simplest one. It does introduce two new parameters, however, where we might have preferred to start with only one, to see how far that goes. As an alternative, we will here explore an exponential of a power of b

$$\mathcal{O}(b) \propto \exp \{-b^d\} \quad (18)$$

where $d \neq 2$ gives deviations from a Gaussian behaviour. We will use the shorthand $\text{ExpOfPow}(d = \dots)$ for such distributions. Note that we do not present an ansatz for $\rho(r)$ from which the $\mathcal{O}(b)$ is derived: in the general case the convolution of two ρ is nontrivial. A peaking of $\mathcal{O}(b)$ at $b = 0$ is related to one of $\rho(r)$ at $r = 0$, however.

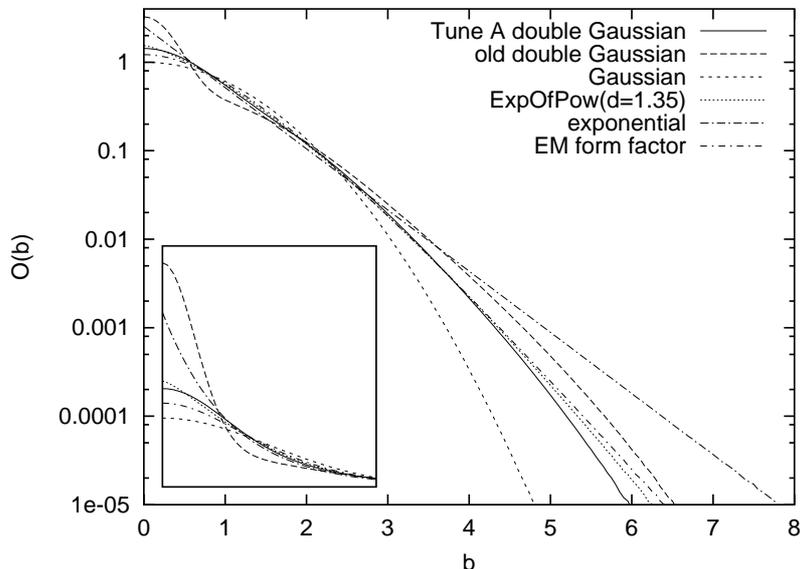


Figure 3: Overlap profile $\mathcal{O}(b)$ for a few different choices. Somewhat arbitrarily the different parametrizations have been normalized to the same area and average b , i.e. same $\int \mathcal{O}(b) d^2b$ and $\int b\mathcal{O}(b) d^2b$. (Recall that we have not specified b in terms of any absolute units, so both a vertical and a horizontal scale factor have to be fixed for each distribution separately.) Insert shows the region $b < 2$ on a linear scale.

A lower d corresponds to an overlap distribution more spiked at $b = 0$ and with a higher tail at large b , Fig. 3, i.e. leads to larger fluctuations. Specifically, the height of the $b = 0$ peak is related to the possibility of having fluctuations out to high multiplicities. To give some feeling, an exponential, $\text{ExpOfPow}(d = 1)$, is not too dissimilar to the old PYTHIA double Gaussian, with $\beta = 0.5$ and $a_2/a_1 = 0.2$. Conveniently, the Tune A double Gaussian, still with $\beta = 0.5$ but now $a_2/a_1 = 0.4$, is well approximated in shape by an $\text{ExpOfPow}(d = 1.35)$. Another alternative, commonly used, is to assume the matter distribution to coincide with the charge distribution, as gauged by the electric form factor $G_E(p_\perp^2) = (1 + p_\perp^2/\mu^2)^{-2}$, with $\mu^2 = 0.71 \text{ GeV}^2$. This gives an $\mathcal{O}(b) \propto (\mu b)^3 K_3(\mu b)$, which also is close in form to Tune A, although somewhat less peaked at small b .

As indicated above, there are two key consequences of a an overlap profile choice. One is the interaction multiplicity distribution and the other the parton-level pedestal effect. These two are illustrated in Figs. 4 and 5, respectively, for $p\bar{p}$ at 1.8 TeV, with $p_{\perp 0} = 2.0 \text{ GeV}$ as in Tune A. The three frames of each

figure illustrate how momentum conservation effects suppress the probability to have an event with large multiplicity. This effect is even stronger now that each

interaction is allowed to undergo full shower evolution, so that it carries away more of the available energy. In the figures, the default lower shower cut-off of 1 GeV has been used; obviously a larger cut-off would give results intermediate to the two lower frames. Further, the possibility of two hard-scattering partons being part of the same shower is not included. Note that the suppression of the high-multiplicity tail implies that a distribution with large fluctuations in reality will have fewer interactions on the average than a less-fluctuating one, if they (as here) start with the same assumed average before the momentum conservation effects are considered. This means that the choice of $p_{\perp 0}$ is somewhat dependent on the one of overlap profile.

Let us now study the hadron-level multiplicity distribution, and begin with UA5 data at 200 and 900 GeV [48]. Tune A then does impressively well, Fig. 6, in spite of primarily having been tuned to pedestal effects rather than multiplicity distributions. In this comparison, we do not put too much emphasis on the low-multiplicity end, which is largely probing diffractive physics. Here the PYTHIA description is known to be too simple, with one or two strings stretched at low p_{\perp} and no hard interactions at all. More relevant is the mismatch in peak position, which mainly is related to the multiplicity in events with only one interaction. Assuming that most hadronization parameters are fixed by e^+e^- data, it is not simple to tune this position. The beam remnant structure does offer some leeway, but actually the defaults are already set towards the end of the sensible range that produces the lower peak position, and still it comes out on the high side.

However, the main impression is of a very good description of the fluctuations to higher multiplicities, better than obtained with the old parameters explored in [1]. Of course, many aspects have changed significantly since then, such as the shape of parton densities at small x . One main difference is that Tune A 90% of the time picks subsequent interactions to be of the $gg \rightarrow gg$ type with colour flow chosen to minimize the string length. Since each further interaction thereby contributes less additional multiplicity, the mean number of interactions can be increased, and this obviates the need for the more extreme double-Gaussian default parameters.

If, nevertheless, one should attempt to modify the Tune A parameters, deviating from its $\text{ExpOfPow}(d = 1.35)$ near equivalent, it would be towards a smaller d , i.e. a slight enhancement of the tail towards high multiplicities. An example is shown in Fig. 6, with $\text{ExpOfPow}(d = 1.2)$ and $p_{\perp 0} = 1.9$ GeV (at 1800 GeV, with the Tune A energy rescaling).

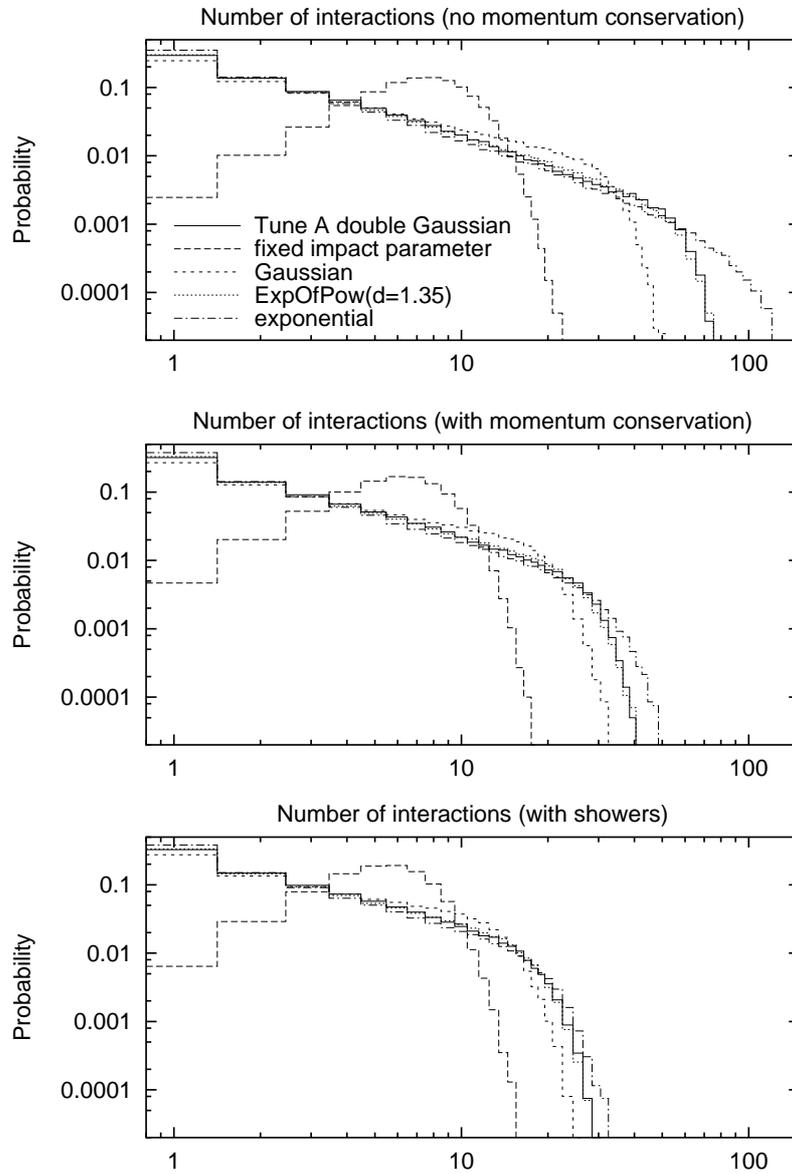


Figure 4: Distribution of the number of interactions for different overlap profiles $\mathcal{O}(b)$, for $p\bar{p}$ at 1.8 TeV, top without momentum conservation constraints, middle with such constraints included but without (initial-state) showers, and bottom also with shower effects included.

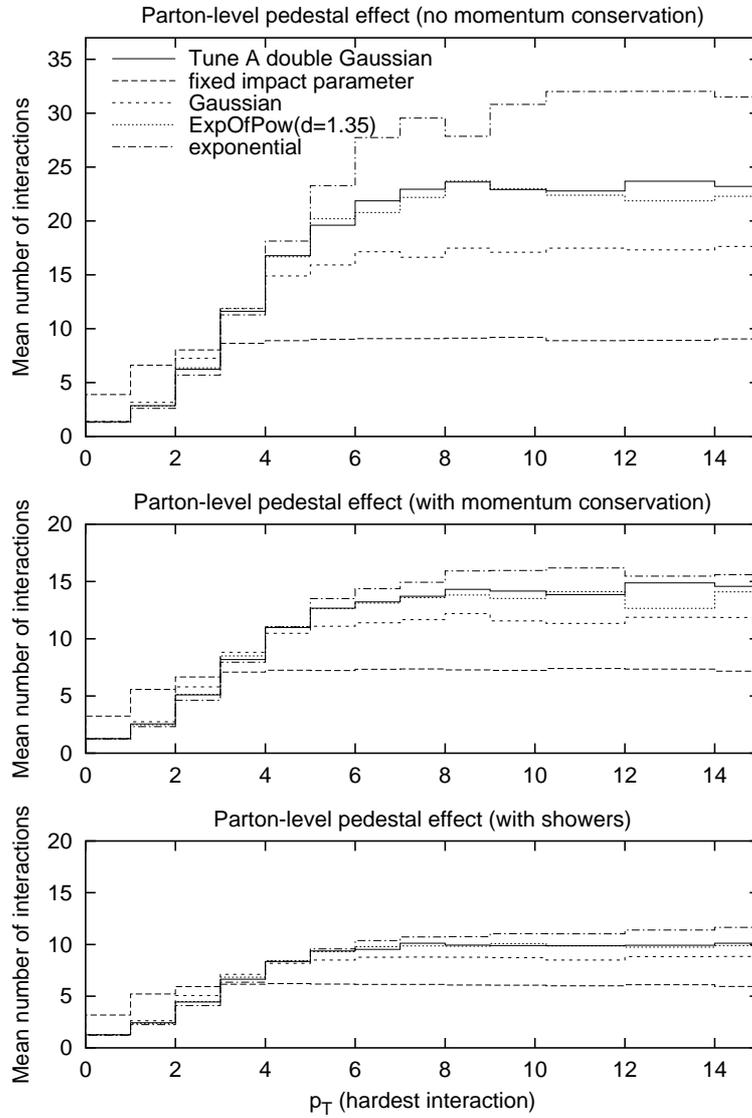


Figure 5: Average number of interactions as a function of the p_{\perp} of the hardest interaction, for $p\bar{p}$ at 1.8 TeV, top without momentum conservation constraints, middle with such constraints included but without (initial-state) showers, and bottom also with shower effects included.

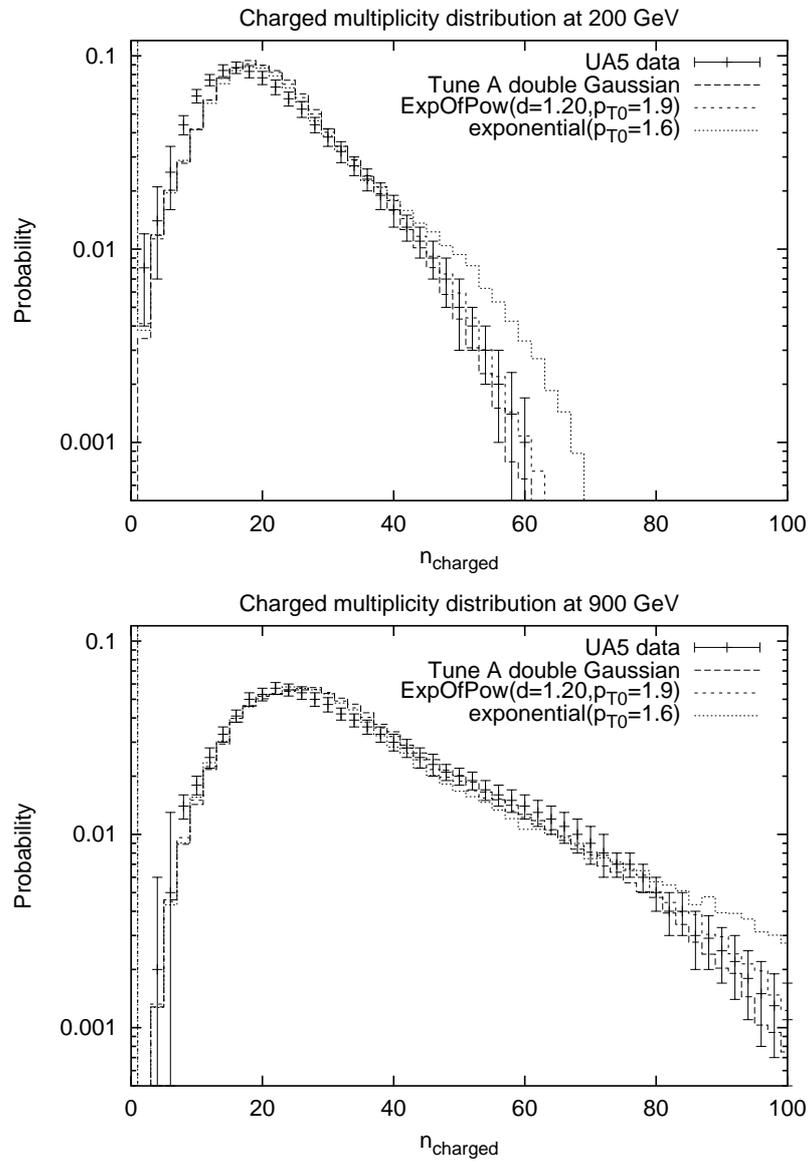


Figure 6: Charged multiplicity distribution at 200 and 900 GeV; different overlap profiles compared with UA5 data [48].

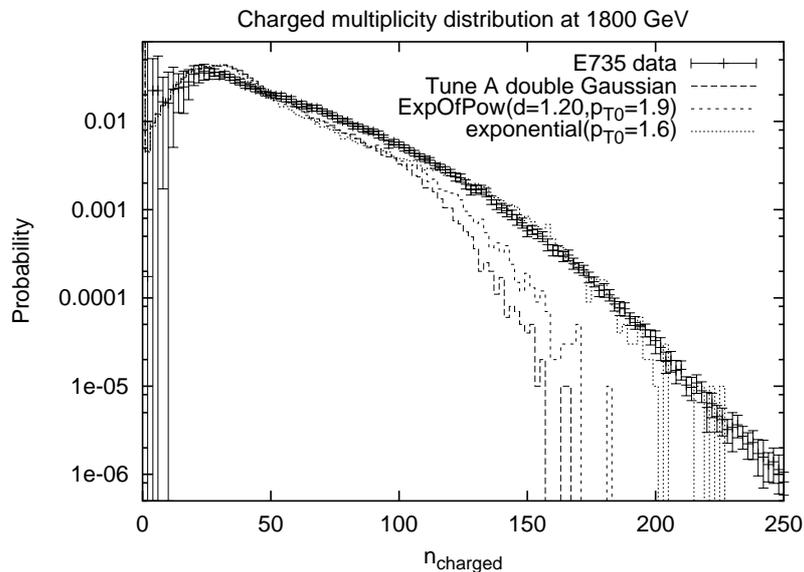


Figure 7: Charged multiplicity distribution at 1800 GeV; different overlap profiles compared with E735 data [49].

However, the nice picture is shattered if one instead considers the E735 data at 1800 GeV [49], Fig. 7. Tune A gives a way too small tail out to large multiplicities, and also the $\text{ExpOfPow}(d = 1.2)$ falls below the data. One would need something like an exponential with a rather low $p_{\perp 0} = 1.6$ GeV to come near the E735 data, and that then disagrees with the lower-energy UA5 data, Fig. 6. The agreement could be improved, but not to the level of Tune A, by playing with the energy dependence of $p_{\perp 0}$. However, the E735 collaboration itself notes that results from the two collaborations are incompatible over the whole UA5 energy range and especially at 546 GeV, where both have data [49]. Furthermore, we do not have the expertise to fully simulate E735 selection criteria, nor to assess the impact of the large acceptance corrections. E735 only covered the pseudorapidity range $|\eta| < 3.25$, so about half of the multiplicity is obtained by extrapolation from the measured region for the 1800 GeV data. UA5 extended further and observed 70%–80%, depending on energy, of its multiplicity.

Obviously new experimental studies would be required to resolve the UA5–E735 ambiguity. As it stands, presumably a tune adjusted to fit E735 would give disagreement with the CDF data that went into Tune A. In this particular case we suspect the differences to be of an experimental origin, but in other cases it could well be that the PYTHIA model is incapable of fitting different

(correct) distributions simultaneously, the model not being perfect. Indeed, speaking in general terms, that is a main reason why we try to improve the model in this article. In this particular case and for the moment being, however, we choose to use the UA5-compatible Tune A as a convenient reference for a realistic multiplicity distribution at Tevatron energies.

4 Correlations in Momentum and Flavour

Consider a hadron undergoing multiple interactions in a collision. Such an object should be described by multi-parton densities, giving the joint probability of simultaneously finding n partons with flavours f_1, \dots, f_n , carrying momentum fractions x_1, \dots, x_n inside the hadron, when probed by interactions at scales Q_1^2, \dots, Q_n^2 . However, just like the standard one-particle-inclusive parton densities, such distributions would involve nonperturbative initial conditions that ultimately would have to be pinned down by experiment. We are nowhere near such a situation: the experimental information on double parton scattering, $n = 2$, boils down to one single number, the σ_{eff} of eq. (10), and for $n \geq 3$ there is no information whatsoever. Wishing to make maximal use of the existing ($n = 1$) information, we thus propose the following strategy.

As described above, the interactions may be generated in an ordered sequence of falling p_\perp . For the hardest interaction, all smaller p_\perp scales may be effectively integrated out of the (unknown) fully correlated distributions, leaving an object described by the standard one-parton distributions, by definition. For the second and subsequent interactions, again all lower- p_\perp scales can be integrated out, but the correlations with the first cannot, and so on.

The general situation is depicted in Fig. 8. This illustrates how, for the i 'th interaction, only the correlations with the $i - 1$ previous interactions need be taken into account, with all lower p_\perp scales integrated out. Note, however, that this is only strictly true for the hard scatterings themselves. The initial-state shower evolution of, say, the first interaction, should exhibit correlations with the i 'th at scales smaller than $p_{\perp i}$. Thus, the p_\perp ordering (or equivalently, a virtuality ordering) is in some sense equivalent to a time ordering, with the harder physics being able to influence the softer physics, but not vice versa. For two interactions of comparable p_\perp this order may appear quite arbitrary, and also should not matter much, but consider the case of one very hard and one very soft interaction. The soft one will then correspond to a long formation time (field regeneration time) [66], $\sim p/p_\perp^2 \sim 1/p_\perp$, and indeed it is to be expected that the hard one can pre-empt or at least modify the soft one, whereas the influence in the other direction would be minor. This gives additional motivation to the choice of a p_\perp ordering of interactions.

The possibility of intertwined shower evolution is not (yet) addressed. Rather, we introduce modified parton densities, that correlate the i 'th interaction and its shower evolution to what happened in the $i - 1$ previous ones, but we

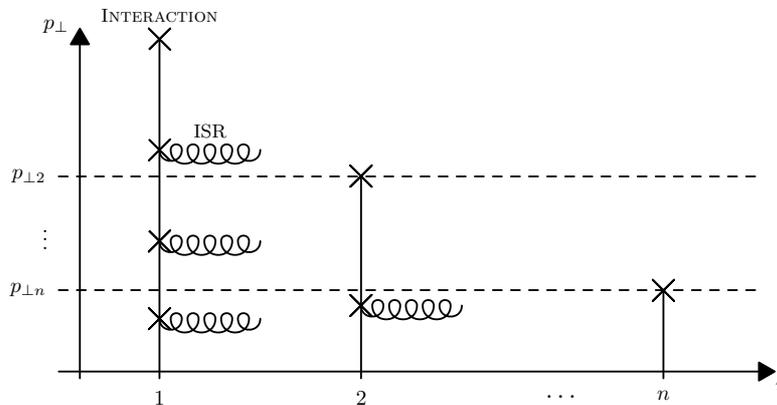


Figure 8: Schematic representation of the evolution of parton shower initiators in a hadron collision with n interactions (see text).

do not let the previous showers be affected by what happens in subsequent interactions. As partons are successively removed from the hadron by hard scatterings at smaller and smaller p_{\perp} scales, the flavour, momentum and colour structure of the remaining object changes. The colour structure in particular is a thorny issue and will be discussed separately, in the next Section. Here, we focus on deriving a set of parton distributions for a hadron after an arbitrary number of interactions have occurred, on the need for assigning a primordial transverse momentum to shower initiators, and on the kinematics of the partons residing in the final beam remnants.

Our general strategy is thus to pick a succession of hard interactions and to associate each interaction with initial- and final-state shower activity, using the parton densities introduced below. The initial-state shower is constructed by backwards evolution [67], back to the shower initiators at some low Q_0 scale, the parton shower cutoff scale. Thus, even if the hard scattering does not involve a valence quark, say, the possibility exists that the shower will reconstruct back to one. This necessitates dealing with quite complicated beam remnant structures. For instance, if two valence quarks have been knocked out of the same baryon in different directions, there will be three quarks, widely separated in momentum space, of which no two may naturally be collapsed to form a diquark system.

In the old model, technical limitations in the way the fragmentation was handled made it impossible to address such remnant systems. Consequently, it was not possible to associate initial-state radiation with the interactions after the first, i.e. the one with the highest p_{\perp} scale, and only a very limited set of $q\bar{q}$ and gg scatterings were allowed.

In a recent article [3], the Lund string model was augmented to include string systems carrying non-zero baryon number, by the introduction of ‘junction fragmentation’. In the context of multiple interactions, this improvement means that almost arbitrarily complicated beam remnants may now be dealt with. Thus, a number of the restrictions that were present in the old model may now be lifted.

4.1 Parton Densities

As mentioned above, we take the standard parton density functions as our starting point in constructing parton distributions for the remnant hadronic object after one or several interactions have occurred. Based on considerations of momentum and flavour conservation we then introduce successive modifications to these distributions.

The first and most trivial observation is that each interaction i removes a momentum fraction x_i from the hadron remnant. This is the fraction carried by the initiator of the initial-state shower, at the shower cutoff scale Q_0 , so that the two initiators of an interaction together carry all the energy and momentum eventually carried by the hard scattering and initial-state shower products combined. To take into account that the total momentum of the remaining object is thereby reduced, already in the old model the parton densities were assumed to scale such that the point $x = 1$ would correspond to the remaining momentum in the beam hadron, rather than the total original beam momentum, cf. eq. (9). In addition to this simple x scaling ansatz we now introduce the possibility of genuine and non-trivial changes in both shape and normalization of the distributions.

Valence Quarks

Whenever a valence quark is knocked out of an incoming hadron, the number of remaining valence quarks of that species should be reduced accordingly. Thus, for a proton, the valence d distribution is completely removed if the valence d quark has been kicked out, whereas the valence u distribution is halved when one of the two is kicked out. In cases where the valence and sea u and d quark distributions are not separately provided from the PDF libraries, we assume that the sea is flavour-antiflavour symmetric, so that one can write e.g.

$$u(x, Q^2) = u_v(x, Q^2) + u_s(x, Q^2) = u_v(x, Q^2) + \bar{u}(x, Q^2). \quad (19)$$

Here and in the following, q_v (q_s) denotes the q valence (sea) distribution. The parametrized u and \bar{u} distributions should then be used to find the relative probability for a kicked-out u quark to be either valence or sea. Explicitly, the quark valence distribution of flavour f after n interactions, $q_{fvn}(x, Q^2)$, is given in terms of the initial distribution, $q_{fv0}(x, Q^2)$, and the ratio of remaining to

original q_f valence quarks, N_{fvn}/N_{fv0} , as:

$$q_{fvn}(x, Q^2) = \frac{N_{fvn}}{N_{fv0}} \frac{1}{X} q_{fv0}\left(\frac{x}{X}, Q^2\right) \quad ; \quad X = 1 - \sum_{i=1}^n x_i, \quad (20)$$

where $N_{uv0} = 2$ and $N_{dv0} = 1$ for the proton, and $x \in [0, X]$ is the fraction of the original beam momentum ($\sum_{i=1}^n x_i$ is the total momentum fraction already taken out of the incoming hadrons by the preceding parton-shower initiators). The Q^2 dependence of q_{fvn} is inherited from the standard parton densities q_{fv0} , and this dependence is reflected both in the choice of a hard scattering and in the backwards evolution. The factor $1/X$ arises since we squeeze the distribution in x while maintaining its area equal to the number of q_f valence quarks originally in the hadron, N_{fv0} , thereby ensuring that the sum rule,

$$\int_0^X q_{fvn}(x, Q^2) dx = N_{fvn}, \quad (21)$$

is respected. There is also the total momentum sum rule,

$$\int_0^X \left(\sum_f q_{fn}(x, Q^2) + g_n(x, Q^2) \right) x dx = X. \quad (22)$$

Without any further change, this sum rule would not be respected since, by removing a valence quark from the parton distributions in the above manner, we also remove a total amount of momentum corresponding to $\langle x_{fv} \rangle$, the average momentum fraction carried by a valence quark of flavour f :

$$\langle x_{fvn}(Q^2) \rangle \equiv \frac{\int_0^X q_{fvn}(x, Q^2) x dx}{\int_0^X q_{fvn}(x, Q^2) dx} = X \langle x_{fv0}(Q^2) \rangle. \quad (23)$$

The removal of $\sum_i x_i$, the total momentum carried by the previously struck partons, has already been taken into account by the ‘squeezing’ in x of the parton distributions (and expressed in eq. (22) by the RHS being equal to X rather than 1). By scaling down the q_v distribution, we are removing an *additional* fraction, $\langle x_{fvn} \rangle$, which must be put back somewhere, in order to maintain the validity of eq. (22).

Strictly speaking, $\langle x_{fv0} \rangle$ of course depends on which specific PDF set is used. Nevertheless, for the purpose at hand this variation is negligible between most modern PDF sets. Hence we make the arbitrary choice of restricting our attention to the values obtained with the CTEQ5L PDF set [6].

More importantly, all the above parton densities depend on the factorization scale Q^2 . This dependence of course carries over to $\langle x_{fv0} \rangle$, for which we assume the functional form

$$\langle x_{fv0}(Q^2) \rangle = \frac{A_f}{1 + B_f \ln\left(\ln(\max(Q^2, 1 \text{ GeV}^2)/\Lambda_{\text{QCD}}^2)\right)}, \quad (24)$$

inspired by the $ds = d \ln(\ln Q^2/\Lambda^2) \propto dQ^2/Q^2 \alpha_s(Q^2)$ pace of evolution, where $d\langle x \rangle/ds \approx -B\langle x \rangle$ suggests a solution of the form $\langle x \rangle \propto \exp(-Bs) \approx 1/(1+Bs)$. Reasonable fits to the CTEQ5L valence quark distributions in the proton are obtained for $A_d = 0.385$, $B_d = 1.60$, $A_u = 0.48$ and $B_u = 1.56$, with the isospin conjugate for neutrons.

Essentially nothing is known about parton densities for other baryons, such as the reasonably long-lived hyperons Λ^0 , $\Sigma^{+,0,-}$, $\Xi^{0,-}$ and Ω^- , which can undergo secondary interactions that one may wish to study. We here use essentially the same parton densities and parameters as for protons. Thus the influence of the larger strange quark mass is neglected, which ought to lead to harder x spectra for s quarks and softer for everything else. The fact that the two proton valence u quarks have a harder distribution than the single d one is carried over to other baryons with two equal quarks, while the average (sum) of the u and d distributions are used for baryons with three unequal (equal) quarks. For mesons and the Vector Meson Dominance part of photons one could use a similar strategy, with π^+ measurements as a starting point instead of protons, while the anomalous part of the photon densities is perturbatively calculable. There are still many further assumptions that would have to go into a complete model of multiple interactions in γp and $\gamma\gamma$ events [12], however, and so far we did not pursue this further.

We now know how much momentum is ‘missing’ in eq. (22). It is not possible to put this momentum back onto valence quarks without changing the shape of the distributions (beyond the mere ‘ x squeezing’) or invalidating eq. (21). Rather, we here assume that the missing momentum is taken up by the sea+gluon distributions, which thus are scaled up slightly when a valence quark is kicked out. This enhancement of the sea+gluon momentum fraction may over- or undercompensate the ‘ x squeezing’ reduction, depending on whether the kicked-out valence quark had a small or large x . However, before the procedure can be discussed in more detail, we must consider another effect which affects the normalization of the sea: changes in the content of the sea itself.

Sea Quarks and their Companions

When a sea quark is kicked out of a hadron, it must leave behind a corresponding antisea parton in the beam remnant, by flavour conservation. We call this a companion quark. In the perturbative approximation the sea quark q_s and its companion q_c come from a gluon branching $g \rightarrow q_s + q_c$, where it is implicitly understood that if q_s is a quark, q_c is its antiquark, and vice versa. This branching often would not be in the perturbative regime, but we choose to make a perturbative ansatz, and also to neglect subsequent perturbative evolution of the q_c distribution. Even if approximate, this procedure should catch the key feature that a sea quark and its companion should not be expected too

far apart in x (or, better, in $\ln x$).

With this approximation, we obtain the q_c distribution from the probability that a sea quark q_s , carrying a momentum fraction x_s , is produced by the branching of a gluon with momentum fraction y , so that the companion has a momentum fraction $x = y - x_s$,

$$\begin{aligned} q_c(x; x_s) &= C \int_0^1 g(y) P_{g \rightarrow q_s q_c}(z) \delta(x_s - zy) dz \\ &= C g(y) P_{g \rightarrow q_s q_c} \left(\frac{x_s}{y} \right) \frac{1}{y} \\ &= C \frac{g(x_s + x)}{x_s + x} P_{g \rightarrow q_s q_c} \left(\frac{x_s}{x_s + x} \right), \end{aligned} \quad (25)$$

with C a normalization constant to be determined below, and $P_{g \rightarrow q_s q_c}$ the DGLAP splitting kernel

$$P_{g \rightarrow q_s q_c}(z) = \frac{1}{2} (z^2 + (1-z)^2). \quad (26)$$

In view of the approximate nature of the procedure, allowing a generic $g(x)$ shape would give disproportionately complex expressions. Instead, the following simple ansatz for the gluon distribution at low Q^2 is used:

$$g(x) \propto \frac{(1-x)^p}{x}, \quad (27)$$

with the integer choices $p = 0, 1, 2, 3, 4$ giving a range of variability for the large- x behaviour of the distribution and the $1/x$ controlling the small- x behaviour. Note that all the above equations are defined assuming no previous energy loss and, as for the valence quarks, should be ‘squeezed’ by a factor X , to ensure momentum conservation.

The overall normalization of a companion quark distribution is obtained by imposing the sum rule:

$$\int_0^{1-x_s} q_c(x; x_s) dx = 1. \quad (28)$$

Inserting eqs. (25)–(27) and inverting, one obtains the normalization constants

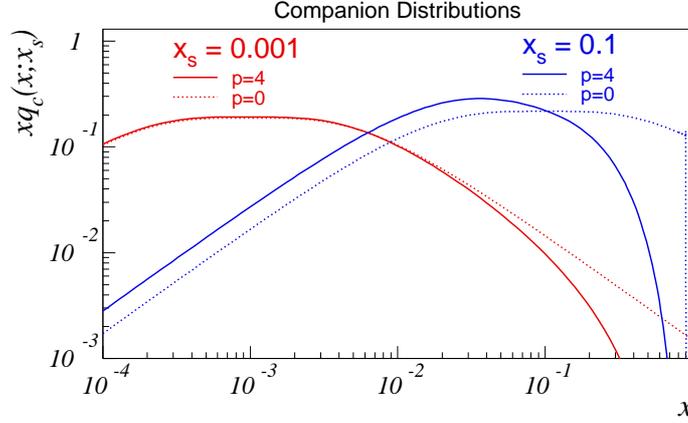


Figure 9: Companion distributions for $p = 4$ (solid lines) and $p = 0$ (dotted lines), for two different values of x_s .

C_p :

$$C_0 = \frac{6x_s}{2 - x_s(3 - 3x_s + 2x_s^2)}, \quad (29)$$

$$C_1 = \frac{6x_s}{2 - x_s^2(3 - x_s) + 3x_s \ln(x_s)}, \quad (30)$$

$$C_2 = \frac{6x_s}{2(1 - x_s)(1 + 4x_s + 4x_s^2) + 6x_s(1 + x_s) \ln(x_s)}, \quad (31)$$

$$C_3 = \frac{12x_s}{4 + x_s(27 - 31x_s^2) + 6x_s(3 + 2x_s(3 + x_s)) \ln(x_s)}, \quad (32)$$

$$C_4 = \frac{6x_s}{2(1 + 2x_s)((1 - x_s)(1 + x_s(10 + x_s)) + 6x_s(1 + x_s) \ln(x_s))}. \quad (33)$$

To illustrate, Fig. 9 shows properly normalized companion momentum distributions for $p = 4$ and $p = 0$, each for two different values of x_s . There are two noteworthy aspects about these distributions. Firstly, for $p = 0$ the discontinuity at the point $x = 1 - x_s$ is merely an artifact of our parametrization of the gluon density, i.e. that $g(x)$ does not vanish at $x = 1$. That problem is absent for more realistic p values; in our continued discussions we will use $p = 4$ as default, this being the closest to the CTEQ5L small- Q^2 gluon distribution. Secondly, the falling gluon distribution convoluted with the almost flat $g \rightarrow q\bar{q}$ splitting kernel give distributions that roughly tend to a constant $q_c(x; x_s) \sim C_p/2x_s^2$ below x_s and exhibit power-like fall-offs $q_c(x; x_s) \propto 1/x^2$ above it, with some modulation of the latter depending on p . In order to display the probability per log x interval, Fig. 9 gives xq_c rather than q_c itself,

and then a peaking occurs around $x \approx x_s$, as should be expected from the symmetric splitting kernel in eq. (26).

Also here, the question arises of ensuring that the total momentum sum rule, eq. (22), is respected, but now the difference has the opposite sign; by adding a companion quark distribution, we are in effect bookkeeping a part of the flavour and momentum content of the sea separately. One possibility is that this momentum comes *only* from the sea+gluons and that the valence quarks are not affected, i.e. that the rest of the sea+gluons fluctuate down, in order to compensate.

The amount of momentum that will have to be compensated for each companion quark, $\langle x_{cn} \rangle = X \langle x_{c0} \rangle$, with $\langle x_{cn} \rangle$ defined in analogy with eq. (23), is straightforward to compute using the distribution in eq. (25) and the normalizations given by eqs. (29)–(33):

$$\langle x_{c0} \rangle_{p=0} = x_s \frac{5 - 9x_s + 6x_s^2 - 2x_s^3 + 3 \ln x_s}{(x_s - 1)(2 - x_s + 2x_s^2)}, \quad (34)$$

$$\langle x_{c0} \rangle_{p=1} = -1 - 3x_s + \frac{2(1 - x_s)^2(1 + x_s + x_s^2)}{2 - 3x_s^2 + x_s^3 + 3x_s \ln x_s}, \quad (35)$$

$$\langle x_{c0} \rangle_{p=2} = \frac{x_s}{4} \frac{19 + 24x_s - 39x_s^2 - 4x_s^3 + 6(1 + 6x_s + 4x_s^2) \ln x_s}{-1 - 3x_s + 3x_s^2 + x_s^3 - 3x_s(1 + x_s) \ln x_s}, \quad (36)$$

$$\langle x_{c0} \rangle_{p=3} = 3x_s \frac{-7 - 21x_s + 15x_s^2 + 13x_s^3 - 2(1 + 9x_s + 12x_s^2 + 2x_s^3) \ln x_s}{4 + 27x_s - 31x_s^3 + 6x_s(3 + 6x_s + 2x_s^2) \ln x_s}, \quad (37)$$

$$\langle x_{c0} \rangle_{p=4} = 3x_s \frac{3(5 + 24x_s - 4x_s^2 - 24x_s^3 - x_s^4) + 4(1 + 12x_s + 24x_s^2 + 8x_s^3) \ln x_s}{8(1 + 2x_s)(-1 - 9x_s + 9x_s^2 + x_s^3 - 6x_s(1 + x_s) \ln x_s)}. \quad (38)$$

Sea Quark and Gluon Density Normalizations

As described above, the normalization of valence and companion distributions is fixed by the respective number of quarks, i.e. the sum rules (for each flavour, f)

$$\int_0^X q_{fvn}(x, Q^2) dx = N_{fvn}, \quad (39)$$

$$\int_0^X q_{fc_j n}(x; x_{s_j}) dx = 1 \quad (\text{for each } j), \quad (40)$$

where X is still the longitudinal momentum fraction left after the n previous interactions and N_{fvn} is the number of q_f valence quarks remaining. The index

j on the companion distribution, q_{fc_jn} , counts different companion quarks of the same flavour, f .

On the other hand, the sea+gluon distributions do not have fixed multiplicities, hence no corresponding sum rules exist for their normalizations. We use this freedom to fulfill the last remaining sum rule, eq. (22), letting the sea+gluon normalizations fluctuate up when we reduce a valence distribution and down when we add a companion distribution. In addition, the requirement of a physical x range is of course maintained by still ‘squeezing’ all distributions into the interval $x \in [0, X]$.

For simplicity, and since eq. (22) only furnishes us with one equation of constraint, we assume the same scale factor for all sea flavours as well as for the gluon, i.e.

$$q_{fs}(x, Q^2) \rightarrow aq_{fs}(x, Q^2), \quad (41)$$

$$g(x, Q^2) \rightarrow ag(x, Q^2). \quad (42)$$

The momentum sum rule now reads:

$$\begin{aligned} 1 &= \frac{1}{X} \int_0^X \left(\sum_f \left[q_{fvn}(x, Q^2) + \sum_j q_{fc_jn}(x; x_j) + aq_{fs}(x, Q^2) \right] + ag_n(x, Q^2) \right) x \, dx \\ &= \int_0^1 \left(\sum_f \left[\frac{N_{fvn}}{N_{fv0}} q_{fv0}(x, Q^2) + \sum_j q_{fc_j0}(x; x_j) + aq_{fs0}(x, Q^2) \right] + ag_0(x, Q^2) \right) x \, dx \\ &= a + \sum_f \int_0^1 \left[\left(\frac{N_{fvn}}{N_{fv0}} - a \right) q_{fv0}(x, Q^2) + \sum_j q_{fc_j0}(x; x_j) \right] x \, dx \\ &= a \left(1 - \sum_f N_{fv0} \langle x_{fv0} \rangle \right) + \sum_f N_{fvn} \langle x_{fv0} \rangle + \sum_{f,j} \langle x_{fc_j0} \rangle, \end{aligned} \quad (43)$$

and hence,

$$a = \frac{1 - \sum_f N_{fvn} \langle x_{fv0} \rangle - \sum_{f,j} \langle x_{fc_j0} \rangle}{1 - \sum_f N_{fv0} \langle x_{fv0} \rangle}. \quad (44)$$

One easily checks that $a = 1$ before the first interaction, as it should be, and that a is driven larger by $N_{fvn} < N_{fv0}$, while introducing companion quarks drives it the opposite way, also as expected.

4.2 Beam Remnants

The longitudinal momenta and flavours of the initiator partons are defined by the sequence of p_\perp -ordered hard scatterings and their associated initial-state showers, as described above. What is left in the beam remnant is then a number of partons, with flavours given by the remaining valence content plus the number of sea quarks required for overall flavour conservation. That

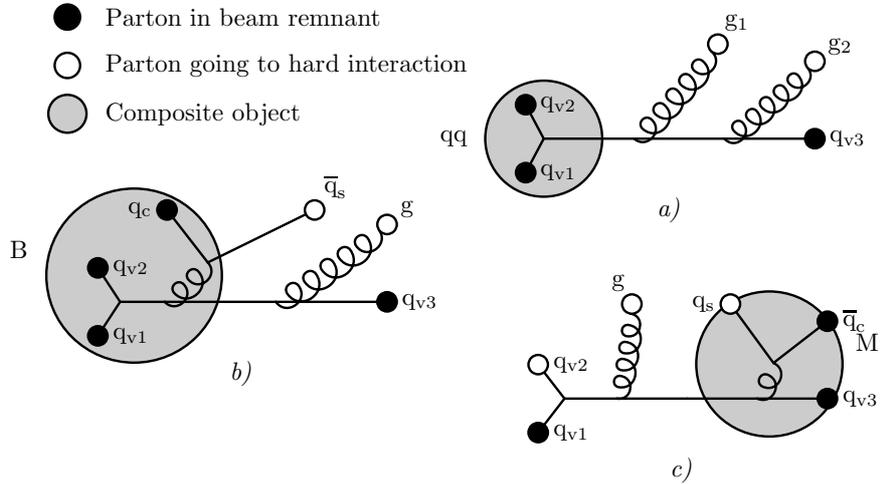


Figure 10: Examples of the formation of composite objects in a baryon beam remnant: (a) diquark, (b) baryon and (c) meson.

is, gluons in the remnant are not explicitly accounted for, but are implicit as confinement clouds around the quarks and as unresolved originators of sea quark pairs.

A remnant may thus contain several objects but, when the colour configuration is studied, simplifications can occur. A colour antitriplet qq pair in the remnant can be associated with a diquark, a colour singlet qqq triplet with a baryon, and a colour singlet $q\bar{q}$ pair with a meson, see Fig. 10. When hadrons are formed, the standard string fragmentation relative probabilities are used to select spin and other quantum numbers, i.e. whether π or ρ , etc.

It is here assumed that the respective pair/triplet has a sufficiently small invariant mass that it can reasonably be projected onto a single composite state. Thus a $q\bar{q}$ system with large invariant mass would define a string that could fragment into several mesons, rather than collapse to a single meson. In principle this could be modeled dynamically, but it would require the introduction of some nonperturbative parameters, to describe the partitioning of the proton into arbitrary-mass subsystems. At this point, we consider it meaningful only to study a few specific scenarios for which partons to allow in the formation of composite objects. We have chosen four such:

1. No composite objects are formed *ab initio*. All partons act as single units, either as endpoints (quarks) or kinks (gluons) on strings that fragment in the normal way.

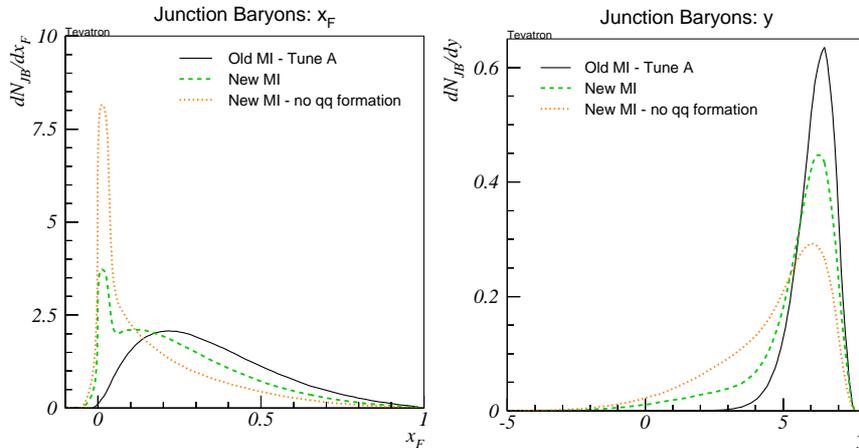


Figure 11: Feynman x (left) and rapidity (right) distributions for junction baryons: distributions are shown for Tune A of the old MI scenario (solid lines), and for the new model with diquark formation in the beam remnant switched on (dashed) and off (dotted).

2. Composite objects may be formed, but only when all partons involved in the formation are valence quarks.
3. The formation of diquarks may involve both valence and sea quarks, but the formation of colour singlet subsystems (i.e. hadrons) is still restricted to involve valence quarks only.
4. Sea quarks may also be used for colour singlet formation.

The idea is thus that (spectator) valence quarks tend to have comparable velocities, while sea quarks can be more spread out and therefore are less likely to form low-mass systems.

Whether composite systems in the beam remnant are formed or not has important consequences for the baryon number flow. For $p\bar{p}$ collisions at 1.8 TeV CM energy, we show in Fig. 11 the Feynman x (left plot) and rapidity (right plot) distributions for the baryon which ‘inherits’ the beam baryon number. We denote this baryon the ‘junction baryon’. To better illustrate what happens to each of the two initial beam baryon numbers separately, only distributions for the junction baryon, not *antibaryon*, are shown. Possibilities 1 and 2 above are compared with the old multiple interactions model (Tune A). One immediately observes that the beam baryon number migrates in a radically different way when diquark formation is allowed or not (compare the dashed and dotted sets of curves). In fact, in the new model it is not possible to reproduce the old

distribution (compare the solid curve). This comes about since, even when all possible diquark formation is allowed in the new model, it is not certain that the beam remnant actually contains the necessary quark content, hence in some fraction of the events the formation of a beam remnant diquark is simply not possible. Here is thus an example where the introduction of more physics into the model has given rise to a qualitatively different expectation: the beam baryon number appears to be stopped to a larger extent than would previously have been expected.

One should note that, also at later stages, a small-mass string piece can collapse to a single hadron, as part of the normal string fragmentation procedure. There, however, it is intended only to cover a rare low-mass tail of systems mainly defined by hard processes and perturbative shower evolution, while the simplifications considered for the above kinematical configurations are quite common and in a nonperturbative context. Beam remnant quarks that were not collapsed in the nonperturbative first stage could in the later stage be collapsed with other partons, e.g. from the showers. Also collapses not allowed in the more restrictive scenarios above could occur at this stage. Whether that happens or not depends on the transverse and longitudinal momenta that will be defined below.

4.3 Primordial k_{\perp}

Until now, we have considered only the longitudinal part of parton momenta. In reality, partons are also expected to have some non-zero k_{\perp} values caused by Fermi motion inside the incoming hadrons. This kind of k_{\perp} is denoted ‘primordial k_{\perp} ’, since it is not generated by the (DGLAP) shower evolution nor from hard interactions, but rather represents an input to the perturbative stages of the event. Based on Fermi motion alone, one would expect values of the order of a few hundred MeV, just like in eq. (6). But to reproduce e.g. the p_{\perp} distributions of Z bosons produced in hadron–hadron collisions, one notes a need for a significantly larger nonperturbative input, either in parton showers or in resummation descriptions [68]. This problem is still awaiting a satisfactory explanation, although some ideas have been explored that might alleviate it [69]. Until such an explanation has been found, we therefore have reason to consider an effective ‘primordial k_{\perp} ’, at the level of the initiators, larger than the one above. For simplicity, a parametrized Q -dependent width

$$\sigma(Q) = \max\left(\sigma_{\min}, \sigma_{\infty} \frac{1}{1 + Q^{\frac{1}{2}}/Q}\right) \quad (45)$$

is introduced, where σ is the width of the two-dimensional Gaussian distribution of the initiator primordial k_{\perp} (so that $\langle k_{\perp}^2 \rangle = \sigma^2$), Q is the scale of the hard interaction, σ_{∞} is the value asymptotically approached as $Q \rightarrow \infty$,

and $Q_{\frac{1}{2}}$ is the Q scale where $\sigma = \frac{1}{2}\sigma_{\infty}$. A reasonable fit to the few available experimental ‘data points’,

$$\begin{aligned}\sigma(Q \sim 1\text{GeV}) &\approx \sigma_{\min} \approx 0.36\text{GeV} && \text{(fragmentation)} \\ \sigma(Q \sim 5 - 10\text{GeV}) &\approx 0.9\text{GeV} && \text{(EMC) [70]} \\ \sigma(Q \sim M_Z) &\approx 2\text{GeV} && \text{(Tevatron) [68],}\end{aligned}\tag{46}$$

is obtained with the values $\sigma_{\infty} = 2.1$ GeV and $Q_{\frac{1}{2}} = 7$ GeV. The σ_{\min} in eq. (45) represents a minimum broadening, at the level of Fermi motion, which we take to be the standard fragmentation p_{\perp} width. In addition to partons participating in relatively soft interactions, this minimum broadening is also applicable to the remnant partons, which by definition do not participate in hard interactions and hence are not naturally associated with a particular Q scale.

Apart from the selection of each individual k_{\perp} , there is also the requirement that the total k_{\perp} of the beam adds up to zero. The question of how partons recoil off one another in transverse momentum space inside hadrons is so far largely unaddressed in the literature. We imagine a few different possibilities here:

1. The primordial transverse momenta are generated at a stage where the partons have low virtualities and hence large wavefunctions. Moreover, at least for that part which is due to Fermi motion, the dynamics responsible for the generation of primordial k_{\perp} is that of a Fermi gas of partons in equilibrium, where each parton has received its total primordial k_{\perp} through a sequence of many collisions with many different partons. Therefore, one possibility is to let the recoil of one parton be shared uniformly among all other initiator and remnant partons.
2. Since Fermi motion alone appears unable to account for the bulk of primordial k_{\perp} in large- Q interactions, there may exist a mechanism, involving presently unknown dynamics, which ties the generation of this k_{\perp} to the presence of a large virtuality in the interaction, for instance by unresolved/unresummed bremsstrahlung radiation off the initiator parton and/or by the k_{\perp} compensations in the shower being of a not strictly local nature (as happens e.g. in the dipole description of parton showers [36]). Regardless of the exact nature of the mechanism, recoils should in this picture primarily be taken up by initiators and beam remnant partons which are close in colour space.
 - (a) The extreme variant is here to let the recoil of a particular parton be taken up by its nearest colour neighbours only.
 - (b) A perhaps more realistic possibility is to let the compensation happen along a parton chain in colour space, with successive dampening of the compensation along the chain.

These three possibilities are included in the present study. However, the possibilities involving colour chains are complicated by the fact that the colour connections between initiator and remnant partons are very poorly known, since no perturbative information is available. These problems will be discussed in more detail in Section 5 below.

Irrespective of which particular method is used to ensure $\sum_i \vec{k}_{\perp i} = 0$, the question now arises how the kinematics of the initially collinear partons should be reinterpreted to include non-zero k_{\perp} assignments. This may be done either by associating the generation of k_{\perp} with the building up of space-like virtualities among the partons, or by keeping the partons massless while allowing (non-perturbatively small) longitudinal momentum transfers between the beam remnants. (In the latter approximation also heavy quarks are kept massless rather than assigned a spacelike virtuality; when initial-state showers are included no heavy quarks need be assigned to the beam remnants, however.) In the first case, the invariant mass of initiator and beam remnant partons combined in each hadron is maintained equal to the original hadron mass, while in the second the mass can be significantly larger. The difference, however, should be considered mostly technical, since the momentum transfers involved are quantitatively small. For this study, the second option is chosen, since this avoids potential technical problems in dealing with string systems having negative mass squares.

For a specific interaction, consider a pair of massless initiator partons in their rest frame, before k_{\perp} is added:

$$p_{1,2} = \frac{\sqrt{\hat{s}}}{2}(1, 0, 0, \pm 1) \quad ; \quad \hat{s} = x_1 x_2 s . \quad (47)$$

With primordial k_{\perp} included, these momentum vectors should now be recast as

$$p_{1,2} = (\sqrt{p_z^2 + p_{\perp 1,2}^2}, \vec{p}_{\perp 1,2}, \pm p_z), \quad (48)$$

if the system should still be at longitudinal rest. Since we are merely reinterpreting the kinematics of the initial-state partons, the centre-of-mass energy, $\sqrt{\hat{s}}$, of the interaction should be left unchanged. To ensure this, it is simple to solve

$$\hat{s} = \left(\sqrt{p_z^2 + p_{\perp 1}^2} + \sqrt{p_z^2 + p_{\perp 1}^2} \right)^2 - (\vec{p}_{\perp 1} + \vec{p}_{\perp 2})^2, \quad (49)$$

to obtain the required p_z as a function of $\vec{p}_{\perp 1}$, $\vec{p}_{\perp 2}$, and \hat{s} :

$$p_z^2 = \frac{\lambda(\hat{s}_{\perp}, p_{\perp 1}^2, p_{\perp 2}^2)}{4\hat{s}_{\perp}} \quad ; \quad \hat{s}_{\perp} \equiv \hat{s} + (\vec{p}_{\perp 1} + \vec{p}_{\perp 2})^2, \quad (50)$$

with λ the standard Källén function,

$$\lambda(a, b, c) = a^2 + b^2 + c^2 - 2ab - 2bc - 2ac. \quad (51)$$

Naturally, only k_{\perp} assignments which result in $p_z^2 > 0$ are acceptable.

4.4 Beam Remnant Longitudinal Momenta

In addition to flavours and transverse momenta, the beam remnants must also together carry the remaining fraction, approximately X (as defined by eq. (20)) of longitudinal momentum. The sharing is based on the character of the remnant constituents. First a fraction x is defined for each constituent, and then these x fractions are rescaled for overall energy and momentum conservation.

Thus a valence quark receives an x picked at random according to a small- Q^2 valence-like parton density, proportional to $(1-x)^a/\sqrt{x}$, where $a = 2$ for a u quark in a proton and $a = 3.5$ for a d quark. A sea quark must be the companion of one of the initiator quarks, and can have an x picked according to the $q_c(x; x_s)$ distribution introduced above. In the rare case that no valence quarks remain and no sea quarks need be added for flavour conservation, the beam remnant is represented by a gluon, carrying all of the beam remnant longitudinal momentum.

Among composite objects, a diquark would naïvely obtain an x given by the sum of its constituent quarks, while baryons and mesons would receive an x equated with the z value obtainable from a fragmentation function, in this study the Lund symmetric fragmentation function. However, earlier studies on quark–diquark remnants [1] have shown that, within the multiple interactions formalism, it is very difficult to accommodate observed remnant multiplicity distributions if the composite system (the diquark) does not take a much larger fraction than implied by the naïve estimate above. Physically, this could correspond to the momentum carried by a surrounding pion/gluon cloud being larger for a composite object than for a single parton. The possibility of enhancing the x values picked for composite objects is therefore retained in the present study.

Finally, once x values (and primordial k_\perp) have been picked for each of the remnants, an overall rescaling is performed such that the remnants together carry the desired longitudinal momentum. The simplest way to accomplish this would be to fix the normalization of the beam remnant x values on each side separately, by requiring conservation of longitudinal momentum, $\sum_i x_{i\text{BR}} = X$ on each side. Unfortunately, the introduction of non-zero k_\perp values with massless partons and on-shell hadrons rules out such a simple approach, since energy would then not be conserved. Instead, small non-zero lightcone momentum fractions in the direction opposite to the parent hadron direction must be allowed. As already noted in Section 4.3, this procedure should be thought of merely as a technical trick, necessitated by insisting on a description in terms of on-shell partons.

The amount of light-cone momentum removed from the remnant system by each pair of initiators, i , is in the overall cm frame of the event (omitting

subscript i to avoid cluttering the notation),

$$\begin{aligned} w^+ &= E^{\text{cm}} + p_z^{\text{cm}} = \gamma(1 + \beta_z)(E'_1 + E'_2), \\ w^- &= E^{\text{cm}} - p_z^{\text{cm}} = \gamma(1 - \beta_z)(E'_1 + E'_2), \end{aligned} \quad (52)$$

where subscripts 1 and 2 denote the initiator partons of the hard scattering on each side respectively, and the primed frame is chosen as the longitudinal rest frame, defined by eq. (48). Then the boost is only along the z direction,

$$\beta_z = \beta = \frac{x_1 - x_2}{x_1 + x_2}, \quad (53)$$

and from equations (48)–(50)

$$E'_1 + E'_2 = \sqrt{\hat{s}_\perp}. \quad (54)$$

Inserting these results in eq. (52) above yields

$$\begin{aligned} w^+ &= \sqrt{\frac{1 + \beta}{1 - \beta}} \sqrt{\hat{s}_\perp} = \sqrt{\frac{x_1}{x_2}} \sqrt{\hat{s}_\perp}, \\ w^- &= \sqrt{\frac{1 - \beta}{1 + \beta}} \sqrt{\hat{s}_\perp} = \sqrt{\frac{x_2}{x_1}} \sqrt{\hat{s}_\perp}. \end{aligned} \quad (55)$$

For vanishing p_\perp this simplifies to the familiar $\hat{s}_\perp = \hat{s} = x_1 x_2 s$, $w^+ = x_1 \sqrt{s}$ and $w^- = x_2 \sqrt{s}$.

The light-cone momenta remaining for the combined beam remnant system are thus:

$$W_{\text{rem}}^+ = \sqrt{s} - \sum_i w_i^+ = \sqrt{s} - \sum_i \sqrt{\frac{x_{i1}}{x_{i2}}} \sqrt{\hat{s}_{\perp i}} \quad (56)$$

$$W_{\text{rem}}^- = \sqrt{s} - \sum_i w_i^- = \sqrt{s} - \sum_i \sqrt{\frac{x_{i2}}{x_{i1}}} \sqrt{\hat{s}_{\perp i}} \quad (57)$$

$$W_{\text{rem}}^2 = W_{\text{rem}}^+ W_{\text{rem}}^-. \quad (58)$$

In extreme cases, it may happen that the hard interactions have removed so much energy and momentum from the beam remnants that the remnant system nominally becomes space-like, $W_{\text{rem}}^2 < 0$, if large k_\perp values have been assigned. Though not strictly speaking unphysical, such a situation could lead to problems at the fragmentation stage. The requirement $W_{\text{rem}}^2 > 0$ is therefore imposed as an additional constraint when primordial k_\perp values are assigned.

The x values picked for the beam remnant partons are now interpreted as fractions of the light-cone momenta, W_{rem}^+ and W_{rem}^- , of the beam remnant

system, modulo an overall rescaling on each side, to leave room for overall momentum conservation. Using index j to refer to beam remnant partons on side 1 and index k for the ones on side 2, we thus make the identification

$$\begin{aligned} p_j^+ &= \alpha x_j W_{\text{rem}}^+ \implies p_j^- = \frac{m_{\perp j}^2}{p_j^+} \\ p_k^- &= \beta x_k W_{\text{rem}}^- \implies p_k^+ = \frac{m_{\perp k}^2}{p_k^-}, \end{aligned} \quad (59)$$

where $m_{\perp}^2 = m^2 + p_{\perp}^2$ and α and β are global normalizations, to be determined from overall energy and momentum conservation in the beam remnant system:

$$W_{\text{rem}}^+ = \sum_j p_j^+ + \sum_k p_k^+ = \alpha W_{\text{rem}}^+ \sum_j x_j + \frac{1}{\beta W_{\text{rem}}^-} \sum_k \frac{m_{\perp k}^2}{x_k} \quad (60)$$

$$W_{\text{rem}}^- = \sum_j p_j^- + \sum_k p_k^- = \beta W_{\text{rem}}^- \sum_k x_k + \frac{1}{\alpha W_{\text{rem}}^+} \sum_j \frac{m_{\perp j}^2}{x_j}. \quad (61)$$

Equating these expressions with eqns. (56) and (57), one obtains for α and β

$$\alpha = \frac{W_{\text{rem}}^2 + W_1^2 - W_2^2 + \lambda^{1/2}(W_{\text{rem}}^2, W_1^2, W_2^2)}{2W_{\text{rem}}^2 \sum_j x_j}, \quad (62)$$

$$\beta = \frac{W_{\text{rem}}^2 + W_2^2 - W_1^2 + \lambda^{1/2}(W_{\text{rem}}^2, W_1^2, W_2^2)}{2W_{\text{rem}}^2 \sum_k x_k}, \quad (63)$$

with the Källén function given by eq. (51) and W_1 and W_2 the total transverse masses of each of the two beam remnant systems,

$$W_1^2 = \left(\sum_j x_j \right) \left(\sum_j \frac{m_{\perp j}^2}{x_j} \right), \quad W_2^2 = \left(\sum_k x_k \right) \left(\sum_k \frac{m_{\perp k}^2}{x_k} \right). \quad (64)$$

Finally, for physical choices of $x_{j,k}$ and primordial k_{\perp} , the sum of the individual remnant transverse masses must be smaller than that of the total remnant system, $W_1 + W_2 < W_{\text{rem}}$. If this is not the case, new k_{\perp} sets and/or new $x_{j,k}$ values are tried, until a physical set of values is found.

5 Colour Correlations and String Topologies

The formalism described in the previous Sections may be used to obtain a sequence of hard scatterings with associated initial- and final-state showers and ordered by the p_{\perp} of the hard interactions they contain. Kinematics is completely specified for all partons involved in the scatterings, in the associated

showers as well as in the left-behind beam remnants. However, at some point, the time evolution of this system results in inter-parton distance scales larger than about a femtometer, where the perturbative QCD description in terms of partons breaks down and must be supplanted by one of hadrons. Now, if each parton hadronized independently of the rest of the event, the information on kinematics and flavours alone would suffice to pass from the language of partons to that of hadrons, but confinement is precisely *not* a strictly local effect. Rather, it is a statement about a (colour singlet) *system* of partons, and hadronization is one concerning the evolution when colour charges inside such a system have been imparted with large momenta *relative to each other*. It is therefore not meaningful to study the hadronization of a single parton in isolation. Instead, it becomes necessary to consider the interplay between colour charges and to take correlations into account when modeling the hadronization process.

In this Section, we begin by considering the hadronization of systems containing non-zero baryon numbers, but where the colours of all participating partons are known. We will use the planar approximation of QCD as a starting point, with a junction picture for the baryons.

Next, since the hard-scattering and parton-shower histories discussed above do not provide sufficient colour information — specifically information about how the different scattering initiators are correlated in colour is completely absent — we consider several possibilities for the assignment of correlated colours to the parton-shower initiators of the scatterings. Combined with the hadronization model, this allows us to study what is obtained under ‘minimal’ assumptions.

Finally, further issues are whether the original colour arrangement survives all the way to the long-distance hadronization era, and whether nonlinear effects arise in the hadronization process itself. That is, with several partons and string pieces moving out from the collision process, these partons and pieces will largely overlap in space and time. We do not know whether such overlaps can lead to colour rearrangements or nontrivial hadronization effects, e.g. of the Bose–Einstein kind. In principle $e^+e^- \rightarrow W^+W^- \rightarrow q_1\bar{q}_2q_3\bar{q}_4$ offers a clean environment to study such crosstalk, but experimental results are inconclusive [79]. In hadronic collisions, Bose–Einstein studies by UA1 and E735 also give a splintered image [80]: the strength parameter of BE effects drops with increasing particle density, consistent with a picture where a higher multiplicity comes from having several independently hadronizing strings, but the BE radius also increases, which suggests correlations between the strings.

5.1 Hadronization

Taking colour interference into account when modeling the hadronization process could easily become an unmanageable task. One simplification (disregard-

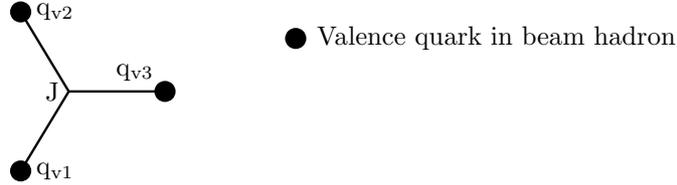


Figure 12: The initial state of a baryon, consisting of 3 valence quarks connected antisymmetrically in colour via a central ‘string junction’, J.

ing baryons for the moment) is to go to the limit with infinitely many colours, $N_C \rightarrow \infty$ [71]. In this limit the confinement force acting on a gluon is twice that on a quark, i.e. the gluon colour and anticolour charges decouple. Further, colour diagrams are planar, so that final-state colour-anticolour pairs are always uniquely matched, via an unbroken colour line through the diagram.

In the Lund string model [4], the two ends of such a colour line define a string piece. The string piece can be viewed as a Lorentz covariant and causal implementation of a linear confinement potential between the two partons. Transverse degrees of freedom here play no dynamical role, but one can visualize the colour field lines as compressed into a tube of a typical hadronic width (~ 1 fm). As the partons move apart and a string piece is stretched out, it can break by the production of new $q\bar{q}$ pairs that screen the endpoint colour charges. The q of one such break and the \bar{q} of an adjacent break together define a meson, which may be unstable and decay further.

The classical example is $e^+e^- \rightarrow q\bar{q}g$, where one may assign a red colour to the q , antired+green to the g and antigreen to the \bar{q} , so that the string consists of two pieces, one $q-g$ and one $g-\bar{q}$. There is no piece directly between q and \bar{q} , with observable consequences in the particle flow [72].

Turning now to baryon beams as the more interesting and difficult example, we picture the initial state of a baryon as consisting of three valence quarks connected in colour via a central ‘junction’, cf. Fig.12. At the most basic level, such a picture finds its motivation by considering the simplest locally gauge invariant operator in $SU(3)$ which carries non-zero baryon number [73]:

$$B_{f_1 f_2 f_3} = \epsilon^{\alpha_1 \alpha_2 \alpha_3} \prod_{i=1}^3 P \left[\exp \left(ig \int_{\mathcal{P}(x, x_i)} G_\mu dx^\mu \right) q_{f_i}(x_i) \right]_{\alpha_i}. \quad (65)$$

Physically, this operator assigns the space-time coordinates x_i to three valence quarks (with flavours f_i and colours α_i) and connects each of them via the gluon field G_μ along the path \mathcal{P} to the point x (with P representing the path-ordering operation), which we may identify as the locus of the string junction. Such ideas

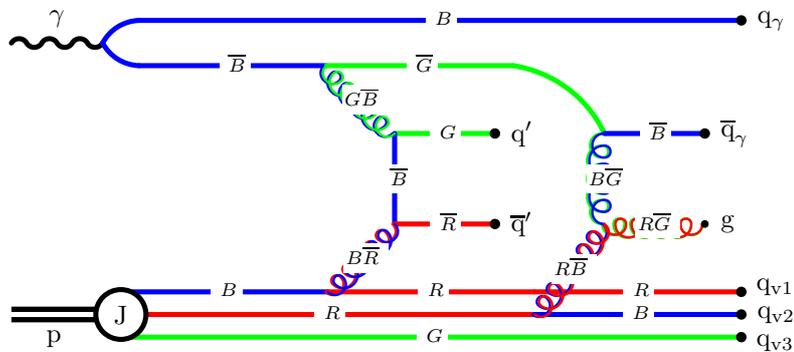


Figure 13: Example of colour assignments in a γp collision with two interactions. Explicit colour labels are shown on each propagator line. In this example, the string system containing the junction is spanned by J connected to q_γ , to q_{v3} , and via g to q' , as can be seen by tracing each of the three colour lines to the junction.

were already introduced in the early string model of hadrons [73–75], and have been used to construct baryon wavefunctions in confinement studies [76]. In a recent article [3], we argued that this picture also arises naturally from string energy minimization considerations.

In a collision, the fact that a gluon carries colour implies that the junction will, in general, be separated in colour space from the original valence quarks. As a simple example, consider a valence $qq \rightarrow qq$ scattering in a pp collision. The exchanged gluon will flip colours, so that each junction becomes attached to a q from the other proton.

Hence the junction may well end up colour-connected to partons — or chains of partons — which are widely separated in momentum space and of which no two may be naturally considered to form a diquark system. To describe the hadronization of such systems, a model capable of addressing colour topologies containing explicit non-zero baryon numbers, here in the form of junctions, becomes necessary. Such a model was first developed in [3], for dealing with the colour topologies that arise in baryon number violating supersymmetric scenarios. In the following, we show how this approach may be applied in a multiple interactions context to describe the physics of beam remnants.

We begin by considering a simplified situation where only one of the initial beam particles is a baryon. Leaving the ambiguities in assigning correlated colours aside for the moment, Fig. 13 gives an example of how the colour structure of a γp collision might look. In this example, the final-state colour-singlet system containing the junction consists of the three string pieces J — q_γ ,

$J-q_{v3}$, and $J-g-q'$. The two other string systems in the event, $q_{v1}-\bar{q}'$ and $q_{v2}-\bar{q}_\gamma$ are standard and do not concern us here.

To understand how the junction system hadronizes, the motion of the junction must first be established. This can be inferred from noting that the opening angle between any pair of the connected string pieces is 120° in the rest frame of the junction, i.e. in that frame the system consisting of the junction and its nearest colour-connected neighbours looks like a perfect Mercedes topology. This is derivable [3] from the action of the classical string [74] (which has a linear potential and thus exerts a constant force), but follows more directly from symmetry arguments.

Note that the junction motion need not be uniform. In the example above, one of the string pieces goes from the junction, via g to q' . At early times, the junction only experiences the pull of its immediate neighbour, g , and the direction of q' is irrelevant. However, as the gluon moves out from the origin, it loses energy to the string traced out behind it. From the point when *all* its energy has been converted to potential energy of the string and this information has propagated back to the junction, it will be the direction of q' which determines the direction of the 'pull' exerted by this string piece on the junction, and not that of the gluon. In the general case, with arbitrarily many gluons, the junction will thus be 'jittering around', being pulled in different directions at different times.

However, rather than trying to trace this jitter in detail — which at any rate is at or below what it is quantum mechanically meaningful to speak about — we choose to define an effective pull of each string on the junction, as if from a single parton with four-momentum given by [3]:

$$p_{\text{pull}} = \sum_{i=1}^n p_i \exp\left(-\sum_{j=1}^{i-1} E_j/E_{\text{norm}}\right) , \quad (66)$$

where the outermost sum runs over the parton chain which defines the string piece, from the junction outwards (in colour space), and where the sum inside the exponent runs over all gluons closer to the junction than the one considered (meaning it vanishes for $i = 1$). The energy normalization parameter E_{norm} is by default associated with the characteristic energy stored in the string at the time of breaking, $E_{\text{norm}} \simeq 1.5\text{GeV}$. Naturally, the energies E_j should be evaluated in the junction rest frame, yet since this is not known to begin with, we use an iterative sequence of successively improved guesses.

With the motion of the junction determined, the fragmentation of the system as a whole may now be addressed. Since the string junction represents a localized topological feature of the gluon/string field, we would not expect the presence of the junction in the string topology to significantly affect the fragmentation in the regions close to the endpoint quarks. Specifically, in an event where each of the three endpoint quarks have large energies in the junction rest frame, the energies of the leading and hence hardest particles of each jet

should agree, on average, with that of an equivalent jet in an ordinary two-jet event.

The hadronization model developed in [3] ensures this by fragmenting each of the string pieces outwards-in, as for a normal $q\bar{q}$ string (in both cases opposite to the physical time ordering of the process). The leading quark of a string piece is combined with a newly created quark-antiquark pair to form a meson plus a new leftover quark, and so on. Parton flavours and hadron spins are selected in a manner identical to that of the ordinary string, as are fragmentation functions and the handling of gluon kinks on the string pieces.

However, junctions were not included in the original string model, so here a new procedure needs to be introduced. If all three string pieces were fragmented in the above way until little energy was left in each to form more hadrons, then it would be extremely unlikely that the resulting leftover system of three unpaired quarks would just happen to have an invariant mass equal to that of any on-shell baryon. While one could in principle amend this by shuffling momentum and energy to other hadrons in the vicinity of the junction, such a procedure would be arbitrary and result in an undesirable and large systematic distortion of the junction baryon spectrum. The way such systematic biases are avoided for ordinary $q\bar{q}$ strings is to alternate between fragmenting the system from the q end and from the \bar{q} end in a random way, so that the hadron pair that is used to ensure overall energy-momentum conservation does not always sit at the same location. Thus, while the distortion is still local in each event, it is smeared out when considering a statistical sample of events.

In the case of a junction system, such a procedure is not immediately applicable. Instead, we first fragment two of the three string pieces, from their respective endpoint quarks inwards. At the point where more energy has been used up for the fragmentation than is available in the piece, the last quark-antiquark pair formed is rejected and the fragmentation is stopped. The two resulting unpaired quarks, one from each fragmented string piece, are then combined into a single diquark, which replaces the junction as the endpoint of the third string piece. Subsequently, this last string piece is fragmented in the normal way, with overall energy and momentum conservation ensured exactly as described for ordinary strings above. In order to minimize the systematics of the distortion and ensure that it is at all possible to produce at least two hadrons from this final string system, we choose to always select the highest energy string piece as the last to be fragmented. It was shown in [3, 77] that this asymmetry in the description does not lead to large systematic effects.

In proton-proton collisions, two junction systems will be present, but it is physically impossible for these to be connected in colour. Hence, the hadronization of each system again proceeds exactly as described above. However, in $p\bar{p}$ collisions a new possibility arises, as depicted in Fig. 14. This simple example goes to illustrate that a junction and an ‘antijunction’ may become colour-connected by the colour exchanges taking place in a given process. In such

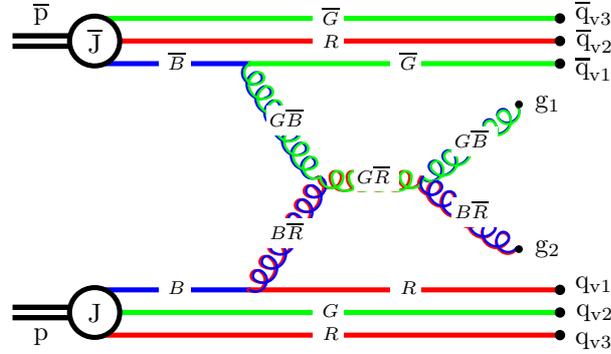


Figure 14: Example of colour assignments in a $p\bar{p}$ collision, with explicit colour labels shown on each propagator line. Note that the blue colour line starting on the junction J is connected via the colour flow of the hard scattering to the antiblue colour line of \bar{J} .

cases, the fragmentation of each of the junctions is no longer disconnected from what happens to the other one; instead the fragmentation of the system as a whole must be considered. The necessary generalization of the principles outlined above to the case of connected junction–junction systems [3] is not very complicated.

As before, two of the three strings from a junction are fragmented first, outwards–in towards the junction, but in this case we always choose these two string pieces to be the ones not connected to the other junction. Diquarks are then formed around each junction exactly as before. What remains is a single string piece, spanned between a diquark at one end and an antidiquark at the other, which can be fragmented in the normal way. In fact, the only truly new question that arises at this point is how to generalize eq. (66) to describe the pull of one junction on another. Here gluons on the string between the two junctions are considered as normally, i.e. their momenta are added, with a suppression factor related to the energy of the intermediate gluons. The partons on the far side of the other junction also contribute their momenta, separately for each of the two strings, with an energy sum suppression now given by the intervening gluons on that particular string, plus the gluons on the junction–junction string.

However, an alternative topology is also possible, where the junction and the antijunction annihilate to produce two separate $q\bar{q}$ systems [3], as illustrated in Fig. 15b. While it is not clear from basic principles how often this should happen, it seems likely that, for a given event, the topology which has the minimal string length is the one selected dynamically. In this case, the string

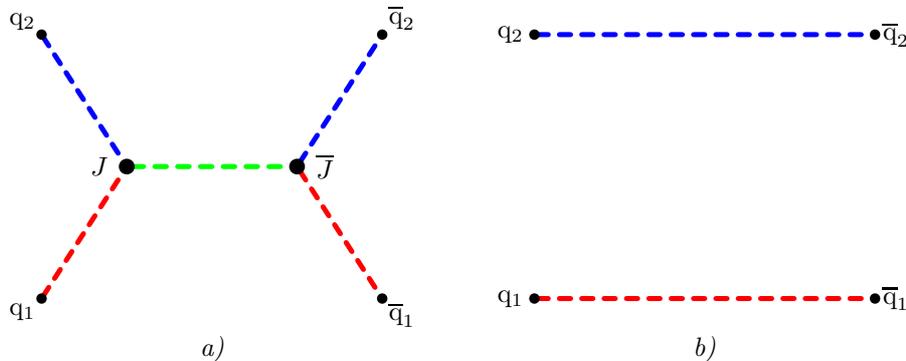


Figure 15: *a)* A string system (dashed lines) spanned between four quarks and containing a junction and an antijunction. *b)* The same parton configuration in colour space but with an alternative string topology. In *a)* the beam baryon numbers will still be present in the final state, while in *b)* they will have disappeared through annihilation.

topology depicted in Fig. 15a would result when the q_1q_2 and $\bar{q}_1\bar{q}_2$ opening angles are small, while the topology in Fig. 15b would result if the $q_1\bar{q}_1$ and $q_2\bar{q}_2$ opening angles are the small ones. Since, in the context we discuss here, the quarks colour-connected to the junction will more often than not reside in the beam remnants, we do not expect annihilation between the incoming baryon numbers to be a large effect. Indeed, for the range of more realistic models that are investigated in Section 6 below, junction–junction annihilation is a feature of less than 1% of the events at Tevatron energies.

An ugly situation occurs in the rare events when the two junctions are connected by *two* colour lines. If these lines contain intermediate gluons, it would be possible but difficult to fragment the system, in particular when the energy of these gluons becomes small. Without intermediate gluons, a first guess would be that the junction and antijunction annihilate to give a simple string spanned between a quark and an antiquark endpoint, so that the original baryon numbers are lost. However, this assumes that the system starts out from a point in space and time, a commonly used approximation in the string language. Viewed in the transverse plane of the collision, the original positions of the junctions and of the hard scatterings involved could well be separated by distances up to a fm, i.e. the intervening strings could have energies up to a GeV. It may then be that the strings can break before the junctions annihilate, so that a baryon–antibaryon pair nevertheless is produced. A detailed modeling

would be required, beyond the scope of the current study, and possibly beyond the validity of the string framework, so for now we choose to reject these rare events.

5.2 Initial-State Colour Correlations

In the planar approximation, $N_C \rightarrow \infty$, a $2 \rightarrow 2$ process, such as $gg \rightarrow gg$, can receive contributions from several possible colour flows, but the cross section for a colour flow is uniquely defined, so that each flow can be selected according to its appropriate weight [78]. Furthermore, in our leading-order parton showers, the colour flow is well-defined in each branching. Within each separate interaction and its associated shower activity, the colour flow can thus be selected unambiguously. In events with only one hard interaction, the colour of the shower initiator can also uniquely be hooked up to that of the beam remnant. Thus a knocked-out quark leaves behind a colour antitriplet beam remnant, a gluon leaves a colour octet beam remnant, and so forth.

Unfortunately, once several interactions are allowed, there is no longer a unique answer how to hook up the different shower initiators with the beam-remnant partons. To illustrate, consider an incident meson out of which n gluons are kicked out. These gluons may be ordered in colour sequence such that the colour of the quark matches the anticolour of one gluon, which then has another colour that matches the anticolour of another gluon, and so on till the antiquark. Obviously there are $n!$ such possible arrangement of the n gluons, each leading to a unique colour topology for the hadronizing partonic system. Perturbation theory has nothing to say about the relative probability for each of these configurations; the colour correlations we now consider arise at scales below the parton-shower cutoff $Q_0 \sim 1$ GeV. Further arrangements would exist if we allowed some of the above gluons to form a separate colour singlet, disconnected from the sequence between the quark and antiquark ends; some of these could contribute to diffractive topologies.

So far we only considered the planar colour topologies of the $N_C \rightarrow \infty$ limit. The real-world $N_C = 3$ offers further complications. First, interference terms of order $1/N_C^2$ — modulated by kinematical factors — arise between different possible colour flows in hard processes. The more partons are involved, the more assignments need be made, and the larger the total uncertainty. Even at the perturbative level it is thus no longer possible to speak of a unique colour arrangement. Second, the situation is more complicated for baryon beams. As described above, the initial state of a baryon, before any scatterings occur, is represented by three valence quarks, connected in a Y-shaped topology via a central junction which acts as a switchyard for the colour flow and which carries the net baryon number. This situation is illustrated in Fig. 12. Each of the gluons considered in the meson-beam example above may now be arranged in colour on either of the three string pieces, leading to a further multiplication

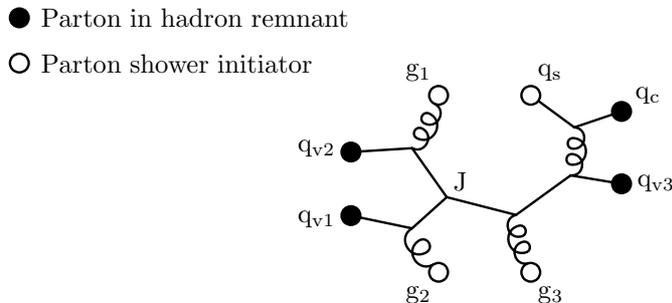


Figure 16: Example of how a given set of parton shower initiators could have been radiated off the initial baryon valence configuration, in the case of the ‘purely random’ correlations discussed in the text. In this example, the baryon number is disconnected from the beam remnant. Instead, it is the final-state partons connected to the colour lines of g_1 , g_2 , and g_3 which determine how the junction moves and hence how the baryon number flows in the event.

of possibilities.

We choose to address this question by determining a sequence of fictitious gluon emissions by which this configuration evolves (in colour space) to give rise to the parton shower initiators and beam remnant partons actually present in a given event. We here assume that only the minimal number of emissions required to obtain the given set of initiators and remnants is dynamically relevant. Further, since sea quarks together with their companion partners can pairwise be associated with a gluon branching below the parton shower cutoff, only gluon emissions remain to be considered. (This also means that a sea quark, in our model, can never form a colour singlet system together with its own companion.)

Random Colour Correlations

The simplest solution would be to assume that Nature arranges these correlations randomly, i.e. that gluons should be attached to the initial quark lines in a random order, see Fig. 16. In this case, the junction (and hence the baryon number) would rarely be colour connected directly to two valence quarks in the beam remnant, even in the quite common case that two such quarks are actually present (multiple valence quark interactions are rare). It should be clear that the migration of the baryon number depends sensitively upon which partons in the final state the junction ends up being connected to (see further Section 5.1). The conclusion is that if the connections are *purely* random, as above, then the baryon number will in general be disconnected from the beam

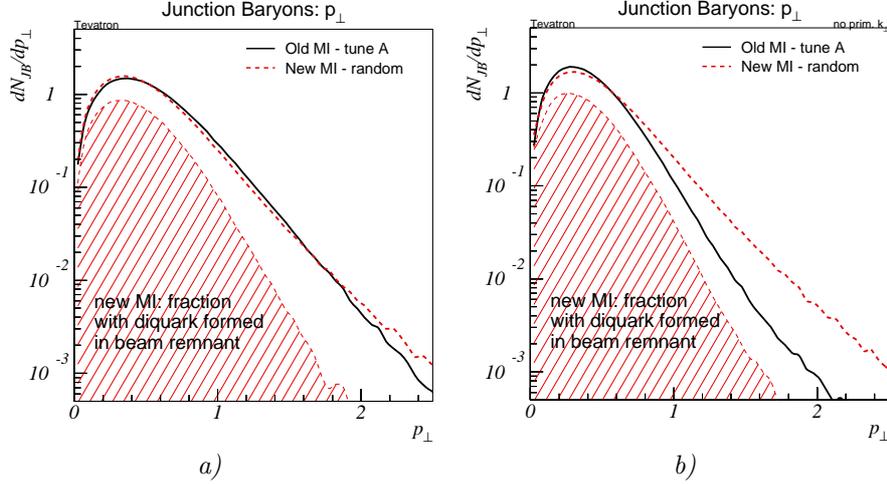


Figure 17: p_{\perp} spectra for junction baryons, *a)* with primordial k_{\perp} switched on, and *b)* with primordial k_{\perp} switched off. The shaded area represents the distribution in the new model of those junction baryons which arose by first forming a diquark in the beam remnant, cf. Section 4.2.

remnant valence quarks. Hence the formation of a diquark in the beam remnant would be rare and the baryon number of the initial state should quite often be able to migrate to small x_F values, as previously illustrated in Fig. 11. One could expect this longitudinal migration to be accompanied by a migration in the transverse plane, such that the junction baryon should generally migrate to larger p_{\perp} values when the junction is allowed to ‘float’ more. However, as Fig. 17a illustrates, no large differences in the total p_{\perp} spectrum are apparent when comparing the new model (thick dashed) with the old Tune A (solid).

The reason that no large transverse migration effect is visible, relative to the old model, is that the latter also has a broader leading-baryon p_{\perp} -spectrum than for normal baryons, as follows. In the old model, the beam remnant diquark (around which the junction baryon forms in the fragmentation) always receives the full primordial k_{\perp} kick from the hard interaction initiator, by default a Gaussian distribution with a width of 1GeV. In the fragmentation process, when the baryon acquires a fraction of the diquark longitudinal momentum, it obtains the same transverse momentum fraction. Additional p_{\perp} will be imparted to the junction baryon from the newly created quark, at the level of 0.36GeV, but with some dependence on the momentum-space location of its nearest neighbour in colour space. Essentially, these two effects combine to yield the solid curve in Fig. 17a.

In the new model, we must distinguish between the old-model-like case

when a diquark is formed in the beam remnant, on one hand, and those cases where the junction is ‘free’ to migrate, on the other.

If a diquark is formed, then it consists of two undisturbed beam remnant quarks which are colour connected to the junction, and the situation is indeed very similar to the old model. The baryon forming around this diquark receives p_{\perp} from three sources. Firstly, the diquark will have some intrinsic primordial k_{\perp} , distributed according to eq. (45). Since the diquark resides entirely within the beam remnant, this k_{\perp} will always be at the level of the fragmentation p_{\perp} . Secondly, primordial k_{\perp} will be imparted to the diquark by recoil effects from other beam remnant and initiator partons. In the case that primordial k_{\perp} kicks are compensated for uniformly by all initiator and remnant partons, it is normally impossible for the diquark to acquire more than a fraction of the hardest interaction initiator’s primordial k_{\perp} . Even when k_{\perp} compensation is more local, by straightforward combinatorics, the more initiators present in the event, the smaller the chances that the initiator parton(s) closest in colour to the diquark is associated with a scattering at large Q^2 . Hence, again according to eq. (45), it is apparent that the diquark usually will not receive a very hard primordial k_{\perp} kick. Thus, such a diquark will in general have a smaller total primordial k_{\perp} than a diquark in the old model. As before, the baryon will keep a large fraction of this diquark p_{\perp} in the fragmentation process, as well as obtaining extra fragmentation p_{\perp} . The net result is a softer junction baryon transverse momentum spectrum than in the old model, as can be verified by comparing the asymptotic slope of the shaded area in Fig. 17a with that of the solid curve. This conclusion is further established by the observation that, when primordial k_{\perp} effects are not included, see Fig. 17b, indeed the spectrum of the old model becomes almost identical to that of the shaded region. In addition, it can already here be recognized that the junction baryon must have larger p_{\perp} in those events where a diquark is *not* formed, by comparing the slopes of the full junction baryon spectrum (dashed curves) with those of the shaded regions in either figure. We now study this further.

If a diquark is not formed, then the junction may *a priori* be colour connected to partons going in widely different directions in the transverse plane. Nonetheless, as was described in Section 5.1, the fragmentation occurs in such a way that the junction baryon is always the last, i.e. normally slowest, hadron to be formed in either of the three directions. Hence, while the colour neighbours of the junction may themselves have large transverse momenta, this momentum will in general be taken by the leading hadrons formed in the fragmentation and not by the junction baryon. Unless two of those partons are going in roughly the same direction in φ , the junction baryon itself will still only obtain a fairly small p_{\perp} . The end result is a rather small p_{\perp} enhancement, that is masked by the decreased primordial k_{\perp} , Fig. 17.

Thus, the main difference in the new model is that the beam baryon number can migrate *longitudinally* to a much larger extent than in the old model. Em-

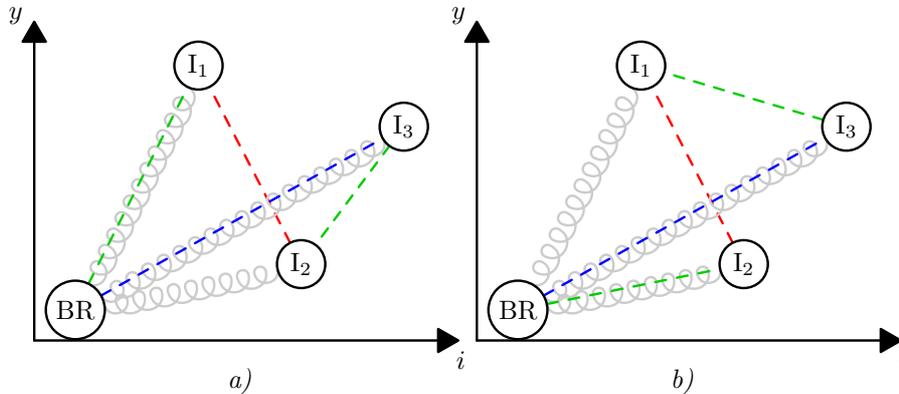


Figure 18: Example of initial-state colour correlations in an imagined event in which three gluons have been knocked out of an incoming hadron by colourless objects (for simplification), with no parton showers. In case *a*) the gluons have been randomly attached in colour to each other and to the beam remnant, as indicated by the dashed lines, whereas in case *b*) the connections follow the rapidities of the hard scattering systems.

pirically, it may be desirable to be able to limit the degree to which this baryon number stopping occurs, and furthermore both perturbative and impact-parameter arguments allow much of the activity to be correlated in ‘hot spot’ regions that leave the rest of the proton largely unaffected. Therefore a free suppression parameter is introduced, such that further gluons more frequently connect to a string piece that has already been disturbed. In this way, gluons would preferentially be found on one of the three colour lines to the junction. This will reduce the amount of baryon number stopping and is an important first modification, but most likely it is not the *only* relevant ordering principle.

Ordered Colour Correlations

With the gluons connected preferentially along one of the three colour lines to the junction, we now address the question of their relative order along that line. If this order is random, then strings will in general be stretched criss-cross in the event. This is illustrated in Fig. 18a for a very simplified situation. However, it is unlikely that such a scenario catches all the relevant physics. More plausible is that, among all the possible final-state colour topologies, those that correspond to the smaller total string length are favoured, all other aspects being the same. One possible way of introducing such correlations is illustrated in Fig. 18b.

In case *a*) of Fig. 18 there are four string pieces criss-crossing the rapidity

range between the systems I_2 and I_3 , while in case b) there are only two string pieces spanning this range. Hence the total string length should on average be smallest for the latter type of correlations. However, the rapidity distance is not the only variable determining the string lengths, also the transverse separations play a role. Moreover, when the interactions exchange colour between the colliding objects, there is no longer a unique correspondence between the colour flow of the initial state and that of the final state.

To investigate these effects quantitatively, we consider three different possibilities for initial-state colour correlations (with the suppression of attachments breaking up the beam remnant applicable to all cases):

1. Random correlations, as in Fig. 18a.
2. Initiator gluons are attached preferentially in those places that order the hard scattering systems in rapidity, as in Fig. 18b. The rapidities are calculated at a stage before primordial k_\perp is included. Hence, $y = \frac{1}{2} \ln \frac{x_1}{x_2}$.

For beam remnant partons, the rapidities are not yet known at the stage discussed here, since the initial-state colour connections are in our framework made *before* primordial k_\perp and beam remnant x values are assigned. However, beam remnant partons are almost by definition characterized by having large longitudinal and small transverse momenta. Thus, we assign a fixed, but otherwise arbitrary, large rapidity to each of the beam remnant partons, in the direction of its parent hadron. Finally, gluons are attached sequentially to the initial valence topology, with the attachments ordered by minimization of the measure

$$\Delta y = |y_g - y_1| + |y_g - y_2|, \quad (67)$$

where y_g is the rapidity associated with the attached gluon and $y_{1,2}$ are the rapidities associated with the partons it is inserted between. For those gluons which appear only as parents of sea quark pairs, the rapidity of the most central of the daughters is used.

Note that, since the same hard-scattering rapidities are used for both beam remnants, the ordering in the two remnants will be closely correlated in this scenario, at least as long as only gluon–gluon interactions are considered.

3. Initiator gluons are attached preferentially in those places that will give rise to the smaller string lengths in the final state. This is the most aggressive possibility, where the actual momentum separations of final-state partons, together with the full colour flow between the two sides of a hadronic collision, is used to determine which gluon attachment will result in the smallest increase in potential energy (string length) of the system, with each gluon being attached one after the other. The measure we use to define the increase

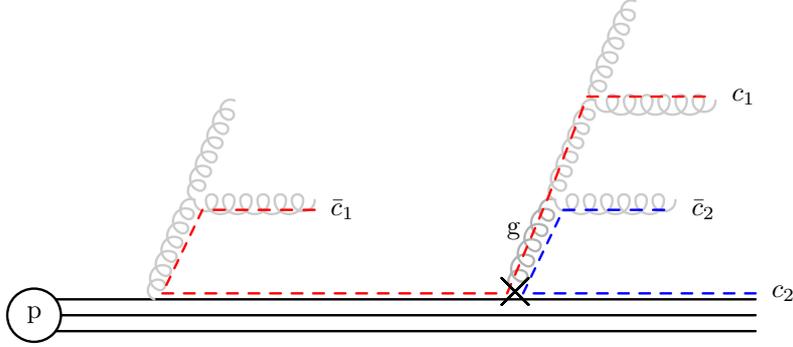


Figure 19: Example showing the colour flow produced by attaching the gluon g at the place indicated by the cross.

in string length, for a particular attachment, is [3, 81]

$$\Delta\lambda = \ln \left[\frac{2}{m_0^2} \frac{(p_{c_1} \cdot p_{\bar{c}_1}) (p_{c_2} \cdot p_{\bar{c}_2})}{(p_{\bar{c}_1} \cdot p_{c_2})} \right], \quad (68)$$

where m_0 is a normalization constant, which drops out when comparing the string lengths of two different gluon attachments, and c_1 (\bar{c}_2) represents the final-state parton carrying the colour (anticolour) index of the attached gluon. To illustrate, Fig. 19 shows the partons that enter the above expression for a specific example. Before the attachment, a single string piece is spanned between the final-state partons that carry the colour indices denoted \bar{c}_1 and c_2 . After the attachment, there are two string pieces, one that is spanned between c_1 and \bar{c}_1 , the other between c_2 and \bar{c}_2 , hence the *increase* in string length is given by the expression eq. (68).

As above, however, note that neither primordial k_\perp nor beam remnant longitudinal momenta have yet been assigned at this stage. Simplified kinematics are therefore set up, to be used only for the purpose of determining the colour connections: the momentum remaining in the beam remnant on each side is divided evenly among the respective remnant partons (junctions are here treated simply as ‘fictitious partons’, receiving the same momentum as the ‘real’ remnant partons), and primordial k_\perp effects are ignored. Thereby, parton pairs involving (at least) two partons in one of the beam remnants will come to have zero invariant mass, hence the total $\Delta\lambda$ will be negative infinity for such pairs. Obviously, this is not desirable; one string piece with vanishing invariant mass should not affect the comparison, hence we impose a minimum invariant mass of m_0 for each string piece. If knowledge of the full kinematics of the final state was available a better choice could of course

be made here. However, these two aspects are intertwined. The kinematics of the final state may depend on the colour connections assumed for the initial state (see Section 4.3 above), and vice versa. Our choice has been to determine the initial-state colour connections first, and then subsequently construct the final-state kinematics, hence some approximation is necessary at this point.

A variable which we have found to be sensitive to the colour connections in an event is the mean p_{\perp} vs. charged multiplicity, $\langle p_{\perp} \rangle(n_{\text{ch}})$ [1]. In scenarios with large string lengths, each additional interaction would result in a large increase in hadron multiplicity. This large multiplicity per interaction means that, in such scenarios, observed average charged multiplicities are reproduced with comparatively large values of $p_{\perp 0}$, i.e. with only a few parton-parton interactions taking place per event. Hence, correspondingly little perturbative p_{\perp} is generated. On the other hand, in scenarios with smaller string lengths, comparatively more interactions would be required to produce the same multiplicity, hence more perturbative p_{\perp} would be generated per charged particle, bringing $\langle p_{\perp} \rangle(n_{\text{ch}})$ up.

In Fig. 20, we show the $\langle p_{\perp} \rangle$ vs. n_{ch} distribution for each of the possibilities described above. In all cases, $p_{\perp 0}$ was first selected so as to give identical average multiplicities, corresponding to the multiplicity obtained with Tune A. Since the $\Delta\lambda$ ordering results in the largest average number of interactions for a given multiplicity, it has the largest average p_{\perp} per particle of the new scenarios, while the random ordering results in the smallest number of interactions.

One also notes that Tune A, which more or less agrees with recent experimental data [56, 58, 61], shows an even steeper rise with n_{ch} than any of the new scenarios. This tune of the old model is such that the partons produced by subsequent scatterings will almost always be hooked up to the existing configuration in the way that minimizes the total string length. This is more or less like the $\Delta\lambda$ ordering described above, but with the essential difference that the $\Delta\lambda$ ordering only concerns the colour lines that are present in the *initial state*, while the ordering of parton attachments in the old model occurs in the *final state*, without any attempt at constructing a consistent colour flow in the event.

From these observations, an interesting inference can be made. By the failure of even the $\Delta\lambda$ ordering of the colour lines in the initial state to describe the $\langle p_{\perp} \rangle(n_{\text{ch}})$ distribution, it appears that the colour flow in physical events cannot be correctly described by merely arranging the colour lines present in the initial state. We imagine two possible causes for this. Firstly, the initial-state showers associated with each scattering are constructed by backwards DGLAP evolution of each scattering initiator separately, down to the shower cutoff scale. This does not take into account the possibility that the showers could be intertwined, i.e. that a parton at low virtuality, but above the shower cutoff scale, could have branched to give rise to *two* higher-virtuality scattering

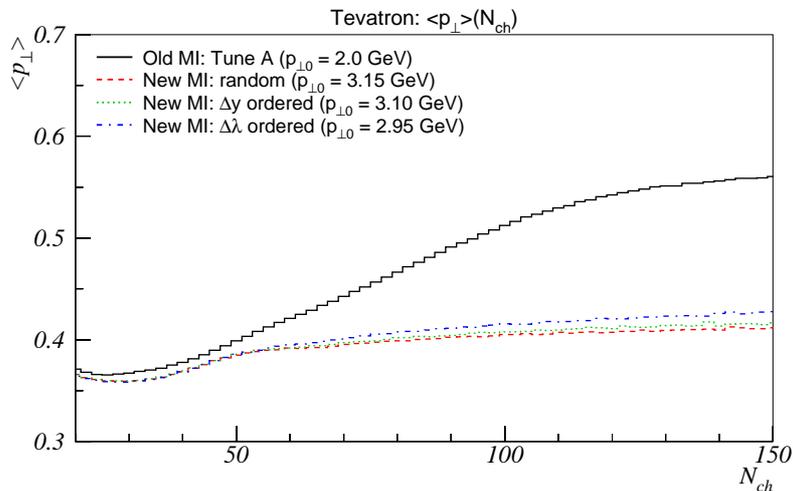


Figure 20: $\langle p_{\perp} \rangle$ vs. n_{ch} at the Tevatron for Tune A (solid lines), and for the new model with random (dashed lines), rapidity ordered (solid lines), and string length ordered (dotted lines) correlations in the initial state. Note that the origo of the plot is *not* at (0,0). For each of the new MI scenarios, $p_{\perp 0}$ was selected to give the same average charged multiplicity as Tune A, with the same impact parameter dependence as Tune A (i.e. a double Gaussian matter distribution).

initiators. Secondly, one or more mechanisms causing colour exchanges between the showers may be active, both in the initial state as well as in the final state. Below we present some first studies related to the topic of colour exchanges, well aware that no simple solutions are to be expected.

5.3 Final-State Colour (Re-)Connections

To investigate how much more we need to ‘mess around with the colours’, we study a crude model of colour exchanges in the final state. Essentially, we rearrange the colour connections between the final-state partons in a manner that, taken to the extreme, will converge on that string configuration which has the smallest total ‘string length’, according to the λ measure introduced above. This corresponds roughly to a minimization of the multiplicity produced when the parton system hadronizes. Note that we only apply this procedure to events where at least two interactions have occurred, to avoid the reconnections leading to large central rapidity gaps, i.e. diffractive topologies.

For a given configuration of final-state partons in momentum and colour

space, we apply an iterative procedure that successively brings down the total string length of the system, by the following steps:

1. First, we have assigned Les Houches Accord style colour tags [82] to all partons, so that each colour tag in the final state is matched by exactly one corresponding anticolour tag in the final state, with one string piece spanned between them. Junctions are special, since colour lines end there. In the following, we do not consider string pieces ending on junctions.
2. Secondly, we decide on a fraction, F , of the colour tags present in the event for which we will attempt to make a reassignment. Note that F can be larger than one, since several different reassignments ($n(n-1)/2$ for n colour tags, neglecting junctions) are normally possible.
3. Next, we select two colour tags at random, c_1 and c_2 . Denoting the final-state parton carrying c_1 colour (anticolour) by i_1 (j_1) and the one carrying c_2 colour (anticolour) by i_2 (j_2), we compute the combined string length, λ , for the two string pieces i_1-j_1 , i_2-j_2 :

$$\lambda = \ln \left(\frac{2p_{i_1} \cdot p_{j_1}}{m_0^2} \frac{2p_{i_2} \cdot p_{j_2}}{m_0^2} \right). \quad (69)$$

By swapping e.g. the anticolours, a different string topology arises, i_1-j_2 , i_2-j_1 , with length

$$\lambda' = \ln \left(\frac{2p_{i_1} \cdot p_{j_2}}{m_0^2} \frac{2p_{i_2} \cdot p_{j_1}}{m_0^2} \right). \quad (70)$$

If $\lambda' < \lambda$, the colour reassignment is accepted, otherwise the original assignments are kept. If the fraction of colour tags tried so far is smaller than F , a new pair of random colour tags is selected.

4. Once the fraction F of colour tags has been tried, two things can happen. If at least one reconnection was made, then the colour topology now looks different, and the entire iteration is restarted. If no reconnection was made, the iteration ends.

Briefly summarized, we thus introduce the fraction F as a free parameter that controls the strength of colour reconnections in the final state.

As stated, this method is very crude and should not be interpreted as representing physics *per se*, but it does allow us to study whether a significant effect can be achieved by manipulating the colour correlations to reduce the string lengths. As illustrated by the dashed histogram in Fig. 21, this is very much the case. Here, we have allowed a large amount of reconnections to occur, $F = 1$, adjusting $p_{\perp 0}$ down so as to reproduce the average charged multiplicity of Tune A. As could be expected with this somewhat extreme choice of param-

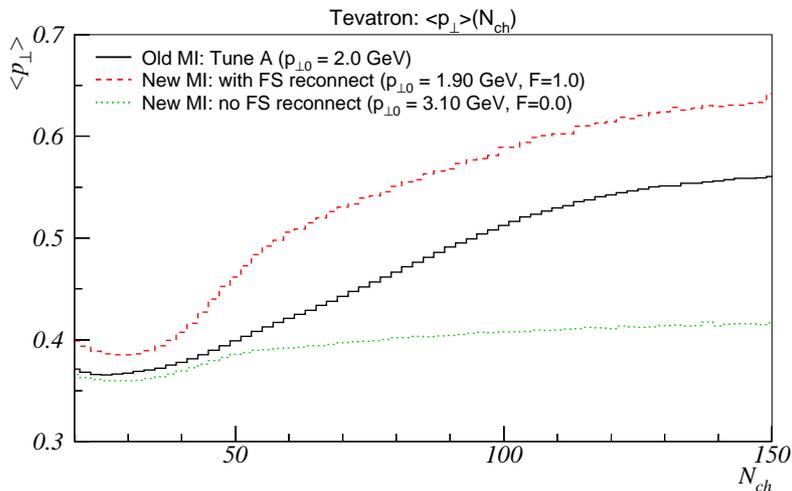


Figure 21: $\langle p_{\perp} \rangle$ vs. n_{ch} at the Tevatron for Tune A (solid line), and for the new model with (dashed line) and without (dotted line) final-state reconnections allowed. Both of the new models use rapidity ordering of the colour lines in the initial state and give the same average charged multiplicity as Tune A, with the same impact parameter dependence as Tune A.

eters, the new model now lies well above the data. (We shall return to more realistic tunings in Section 6.)

However, this result should not be taken as evidence for the existence of colour reconnections in physical events. Rather it allows us to infer that, by changing the colour structure of events, it should be possible to obtain agreement with the data within our framework. It is encouraging that, by studying and attempting to describe this distribution, we may learn interesting lessons concerning the highly non-trivial issue of colour flow in hadronic interactions. We therefore plan to go further, to construct more physically motivated models for colour rearrangement between partons, both in the initial state and in the final state, and also to allow for the possibility of intertwining the initial-state showers.

Although all of these issues appear almost hopelessly complicated from the point of view of pure QCD, the salient features of the resulting physics may not be all that hard to penetrate. For instance, we imagine that the probability for two hard-scattering initiators to have originated from a common branching should be proportional to the probability that their spatial wavefunctions overlap, i.e. two very high-virtuality initial-state partons associated with different scatterings are most likely uncorrelated, since they only resolve

Model name	IS Colour Ordering	$p_{\perp 0}$ [GeV]	F	Tevatron			LHC		
				$\langle n_{\text{int}} \rangle$	$\langle n_{\text{part}} \rangle$	$\langle n_{\text{rcp}} \rangle$	$\langle n_{\text{int}} \rangle$	$\langle n_{\text{part}} \rangle$	$\langle n_{\text{rcp}} \rangle$
Ran	random	2.50	0.55	3.3	21.5	18	4.2	40.2	45
Rap	Δy	2.40	0.55	3.6	22.8	19	4.5	43.5	49
Lam	$\Delta \lambda$	2.30	0.65	3.9	24.5	20	4.8	45.8	52
Tune A	–	2.00	–	5.7	19.2	–	6.9	27.7	–

Table 1: Parameters of the three models investigated in the text, and for Tune A where applicable. Also shown for each model is the average number of parton–parton interactions, $\langle n_{\text{int}} \rangle$, the average number of final-state partons, $\langle n_{\text{part}} \rangle$, and the average number of colour reconnections taking place, $\langle n_{\text{rcp}} \rangle$, per min-bias collision at the Tevatron and at the LHC.

very small distance scales in their parent hadron. On the other hand, two low–virtuality partons, even though associated with different scatterings, may very well have come from one and the same parent parton, since their wavefunctions are comparatively much larger. We will explore such and other ideas in a future study.

6 Model Studies

In this Section we concentrate on illustrating the properties of models that, as a baseline, roughly reproduce the charged multiplicity distribution of Tune A in $p\bar{p}$ collisions at a centre-of-mass energy of 1800GeV. Our studies only concern inelastic nondiffractive events, i.e. essentially the same as the experimentally defined (trigger-dependent) “min-bias” event sample; we will here use the two concepts interchangeably.

A general feature of the new multiple interactions modeling is that the added parton showers and the less efficient string energy minimization result in a higher multiplicity per interaction than Tune A. In order to arrive at the same average hadron multiplicity as Tune A, without too much colour reconnection required in the final state, a generally larger $p_{\perp 0}$ cutoff should be used. Table 1 lists three different tunes of the new framework, with successively smaller $p_{\perp 0}$ values and with different schemes for the initial-state colour correlations.

- The “Ran” model is based on a random ordering of the initial-state colour correlations, with a fairly large suppression of initiator gluon attachments to colour lines wholly within the beam remnant. Since only a minimal ordering of the colour correlations in the initial state are thus imposed, each additional interaction will *ab initio* give rise to a relatively large increase in hadron multiplicity. Therefore, a comparatively large cutoff $p_{\perp 0}$ is used, and the F parameter — controlling the amount of final-state reconnections — is likewise

chosen fairly large, so as to get the correct average charged multiplicity.

- The “Rap” model uses the Δy measure introduced above to order the initial-state colour connections. $p_{\perp 0}$ can thus here be slightly smaller, allowing more interactions on the average (with the same F fraction) for the same average charged multiplicity.
- The “Lam” model employs the $\Delta\lambda$ ordering of the initial-state colour correlations. In principle, this model should provide the most ordered initial state of the three, and thus allow a smaller $p_{\perp 0}$ and/or F . Unfortunately, the earlier-mentioned limitations, that the beam remnant kinematics is not fully fixed when the minimization is performed, leads to final string lengths which are not significantly shorter than for the Δy ordering. Choosing a smaller $p_{\perp 0}$ for this tune, the F fraction is consequently also required to be slightly higher, in order to reproduce the Tune A average charged multiplicity.

Observe that all three models have a significant number of reconnections per event, cf. $\langle n_{\text{rcp}} \rangle$ in Table 1. As a fraction of the total number of potential colour rearrangements it is below the 10% level, but one should keep in mind that several clusters of partons appear in reasonably collimated jets, where reconnections would not be expected anyway. In this perspective, the amount of reconnections is quite significant.

Common for the new models is a rather smooth overlap profile ExpOf-Pow(1.8), as compared to the more peaked double Gaussian of Tune A, this to better reproduce the shape of the Tune A multiplicity distribution. In addition, all of the new models assume that primordial k_{\perp} kicks are compensated uniformly among all other initiator partons, that composite objects can only be formed in the beam remnant by valence quarks, and that initial-state colour connections breaking up the beam remnant are suppressed, so that the relative probability of attaching a gluon between two remnant partons (i.e. breaking up the remnant) as compared to an attachment where at least one ‘leg’ is outside the remnant is 0.01, whenever the latter type of attachment is possible.

Fig. 22 shows the Tevatron multiplicity distributions of these models, as compared to Tune A. It is apparent that, while the average charged multiplicity is the same, the shape of the Tune A multiplicity distribution is not exactly reproduced by any of the models here investigated. This should not be taken too seriously; our aim is not to present full-fledged tunes, rather it is to explore the general properties of the new framework, and how these compare with those of Tune A.

Below, we first present comparisons for $p\bar{p}$ min-bias collisions at 1.8 TeV CM energy, highlighting the differences (and similarities) between the new models and Tune A. Thereafter, we apply the same models to the case of pp min-bias events at 14 TeV CM energy.

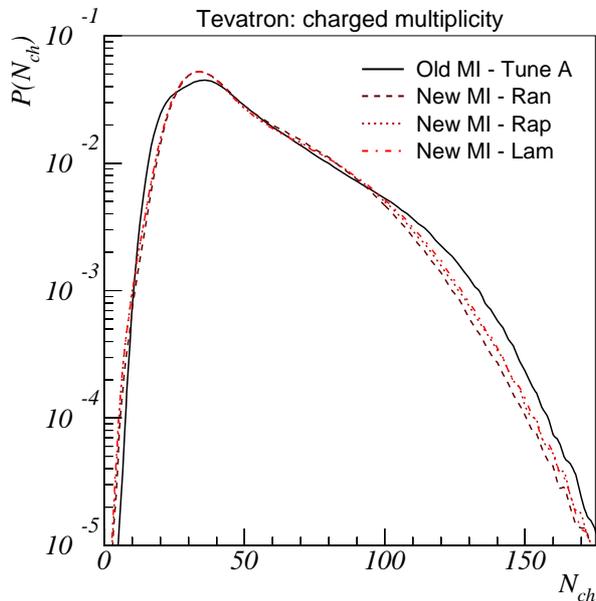


Figure 22: Multiplicity distributions for the Tevatron as obtained with Tune A (solid), and the Ran (dashed), Rap (dotted), and Lam (dash-dotted) models defined in Table 1. In all cases, the average charged multiplicity is 49.5 (within ± 0.5).

6.1 Comparisons at the Tevatron

Despite the differences between the old and new frameworks, the starting point in both cases is still that of a perturbative sequence of p_{\perp} -ordered scatterings. Especially for the hardest partons there should thus be next to no difference between the old and new frameworks. An illustration of this is given in Fig. 23, where the probability of finding a jet with transverse energy E_{\perp} is plotted against E_{\perp} . A simple cone algorithm with cone size $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} = 0.7$ has been used to cluster the jets, and only particles with $|\eta| < 2.5$ are included. As can be observed, there is hardly any difference between Tune A and the new models here. Further, the models exhibit similar charged particle spectra both in transverse momentum and in rapidity, Fig. 24, with a slightly harder p_{\perp} spectrum and, consequently, a slightly more central y one in the new models.

Having thus convinced ourselves that the overall features of the models agree, we turn to the aspects in which the new and old scenarios are expected to differ. One significant change is the possibility to knock out several valence quarks from the beam hadron. To quantify, the number of quarks (excluding

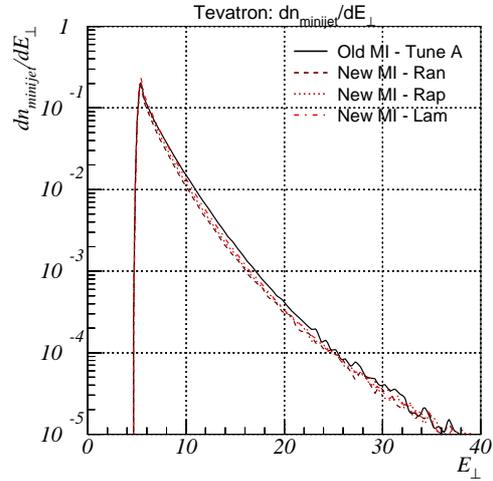


Figure 23: The number of jets as a function of jet E_{\perp} (for $E_{\perp} > 5\text{GeV}$) in min-bias collisions at the Tevatron. Results are shown for Tune A (solid), Ran (dashed), Rap (dotted), and Lam (dash-dotted), as defined in Table 1.

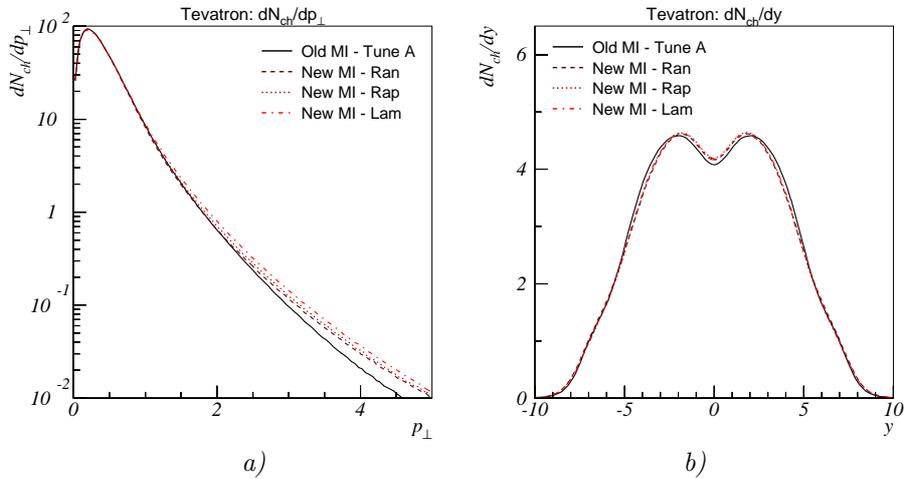


Figure 24: Charged particle p_{\perp} and y spectra at the Tevatron. Results are shown for Tune A (solid), Ran (dashed), Rap (dotted), and Lam (dash-dotted), as defined in Table 1.

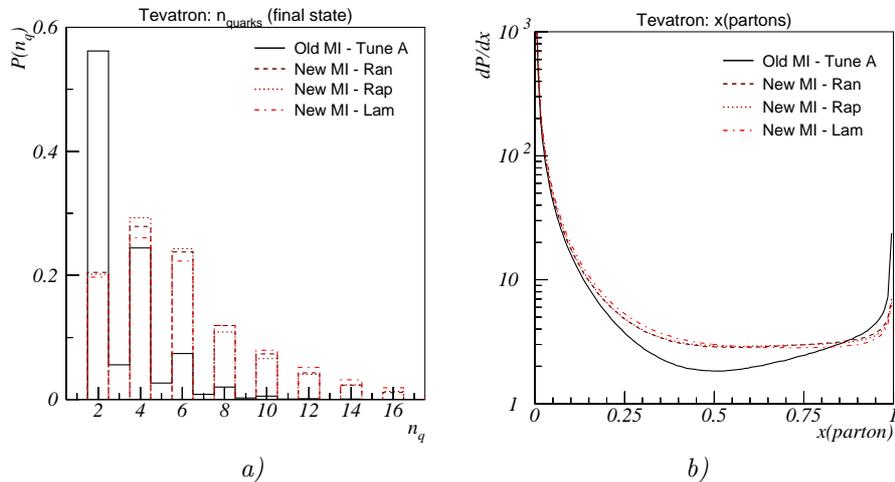


Figure 25: *a)* Number of final-state quarks (not counting diquarks) and *b)* final-state parton x values at the Tevatron. Results are shown for Tune A (solid), Ran (dashed), Rap (dotted), and Lam (dash-dotted), as defined in Table 1.

diquarks) in the final state is illustrated in Fig. 25a. As can readily be observed, final-state quarks are much more abundant in the new scenarios. Valence quark interactions are, however, not the only cause of this. Diquarks are not included in Fig. 25a, and (as discussed in Section 4.2) these are less frequently formed in the new scenarios, hence more of the quark content here appears as individual quarks in the final state. Further illustration of this is given by Fig. 25b, where the $x = 2E/\sqrt{s}$ values of all final-state partons, including diquarks, are shown. The peak towards low values comes mainly from gluons and is the same for Tune A and the new models. However, some of the content of the $x = 1$ peak in Tune A has vanished in the new models and has been replaced by a larger plateau at intermediate x , since a fraction of large- x diquarks has been split up into individual valence-like quarks.

Another point where the models differ is in the treatment of the beam remnant, especially concerning the flow of baryon number. Fig. 26a shows the distribution of baryons minus the distribution of antibaryons as a function of rapidity, illustrating one way of experimentally probing the location of the beam baryon numbers in the final state. As expected, the distribution is more peaked in the old scenario, where the beam baryon number is ‘locked’ inside the escaping remnant diquark. Further illustration of the baryon number stopping is given by Fig. 26b, which shows the rapidity distribution of the final-state baryon carrying the baryon number of the incoming proton. Note especially the long tail in the new models, even extending to negative rapidities. The

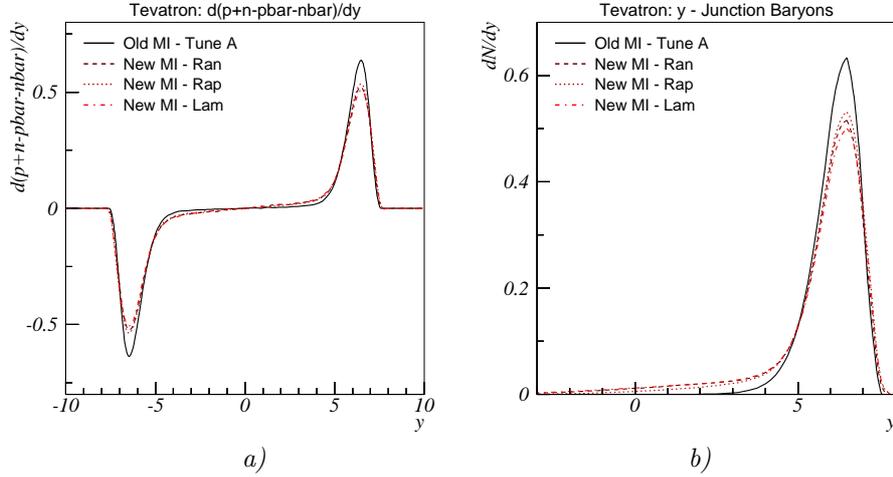


Figure 26: *a)* $n(p) + n(n) - n(\bar{p}) - n(\bar{n})$ and *b)* junction baryon rapidity distributions at the Tevatron. Results are shown for Tune A (solid), Ran (dashed), Rap (dotted), and Lam (dash-dotted), as defined in Table 1.

height of this tail depends on a set of model parameters, such as the choice of initial- and final-state colour correlations and the rules for diquark formation, cf. Fig. 11 for more extreme scenarios.

The fact that the baryon number is not so closely associated with the valence quarks in the new model has an interesting side effect: apart from migrating in physical space, the initial-state baryon number can also ‘migrate’ in flavour to a much larger extent than before. When the junction is resolved, the flavour selection rules of ordinary string fragmentation take over in determining the flavour composition of the junction baryon. Thereby, it becomes possible for the junction baryon to have a larger strangeness number S than before, as shown in Fig. 27. The last bin, $|S| = 3$, is actually empty for Tune A, since in the old model it is impossible to produce an Ω^- from the incoming beam baryon number. We note that heavy-ion experiments do observe a ratio $\Omega^-/\Omega^+ > 1$ [83], which is in contradiction with the old string model but which would be more in line with expectations based on the junction scenario introduced here.

Finally, Fig. 28 shows the $\langle p_\perp \rangle (n_{\text{ch}})$ distributions. As previously noted, Fig. 21, this distribution is very sensitive to the colour correlations present at the hadronization stage. Since this is one of the major open issues remaining, it is not surprising that the agreement here is far from perfect: the new models exhibit a too early rise to a too low plateau, as compared to Tune A. While it is possible to obtain a much better agreement by varying the scheme adopted for the colour reconnections in the final state, our attempts in this direction

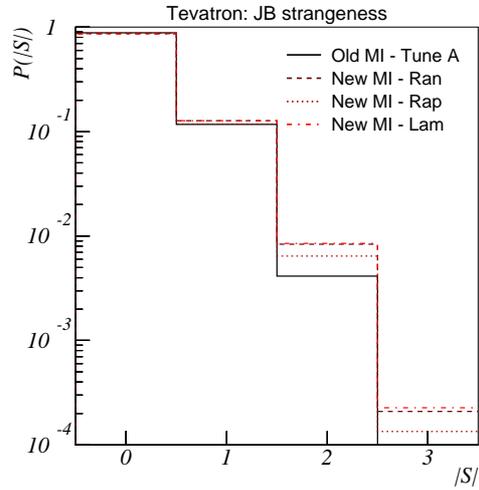


Figure 27: Junction baryon strangeness. Results are shown for Tune A (solid), Ran (dashed), Rap (dotted), and Lam (dash-dotted), as defined in Table 1.

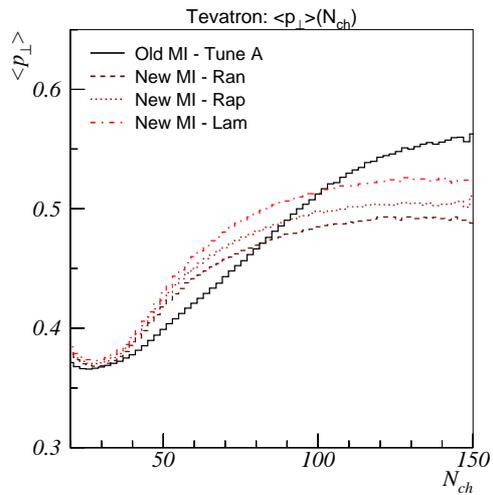


Figure 28: Average p_{\perp} as a function of charged multiplicity for min-bias collisions at the Tevatron. Results are shown for Tune A (solid), Ran (dashed), Rap (dotted), and Lam (dash-dotted), as defined in Table 1.

have so far led to poorer descriptions of other distributions, the charged multiplicity distribution in particular. Moreover, the colour reconnection scheme adopted here is meant only as an instructive example, not as a model of the physics taking place. Our plan is to continue the study of colour correlations in more depth. In the context of such studies, it is encouraging to note that the acute sensitivity of the $\langle p_{\perp} \rangle (n_{\text{ch}})$ distribution to these aspects makes its proper description a prime test for any physical model of the colour flow.

6.2 Comparisons at the LHC

Turning now to the situation at the LHC, the lack of experimental constraints increases the uncertainties. We here focus on only a subset of these, assuming the same energy scaling of the $p_{\perp 0}$ cutoff as for Tune A, proportional to $E_{\text{cm}}^{0.25}$, and using the same parton distributions in all cases. We note that these aspects do constitute important sources of uncertainty in our ability to make trustworthy ‘forecasts’ for the LHC. Here, however, our aim is merely to compare alternative scenarios under identical boundary conditions.

Fig. 29 shows the LHC charged multiplicity distribution for the same models as used for the Tevatron, but with the $p_{\perp 0}$ cutoff scaled to the LHC CM energy. The new models exhibit average multiplicities of 6–8 more charged particles per event than the Tune A value, $\langle n_{\text{ch}} \rangle = 81$. This illustrates a general effect in the new models, which is due to the increased shower activity arising from associating also the sub-leading interactions with initial- and final-state cascades. The larger available phase space at higher energies implies that showers are more important at LHC, cf. the average number of final-state partons, $\langle n_{\text{part}} \rangle$, in Table 1. All else being equal this causes the multiplicity to increase more rapidly with energy than in the old model. (The strange bump on the Tune A distribution at low multiplicities is merely an artifact of the way parton distributions at low Q^2 are handled in that model.)

The addition of parton showers also increases the total amount of partonic transverse energy, but owing to a partial cancellation of the effects of radiation in the initial state (boosting some partons to larger p_{\perp}) and in the final state (jet broadening), the jet rates come out similar, as depicted in Fig. 30 (adopting the same cone algorithm and $|\eta| < 2.5$ region as before).

Fig. 31 compares the junction baryon p_{\perp} distributions at the Tevatron (left plot) and at the LHC (right plot), for Tune A and the new models. An interesting difference is that the junction baryon can be significantly harder in p_{\perp} at the LHC than at the Tevatron in the new models, whereas Tune A exhibits spectra which are almost identical between the two energies. This is due to the intrinsic difference between the way primordial k_{\perp} is treated in the two frameworks. In the old model, the width of the primordial k_{\perp} distribution for the parton initiating the hardest scattering is fixed, to 1GeV by default, hence there is no mechanism that would allow the junction baryon spectrum to depend on

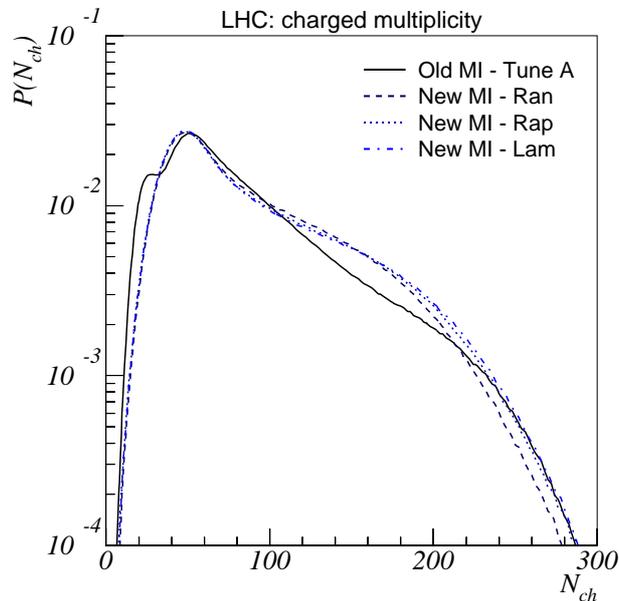


Figure 29: Multiplicity distributions for the LHC as obtained with Tune A (solid), and the Ran (dashed), Rap (dotted), and Lam (dash-dotted) models defined in Table 1. For Tune A, the average charged multiplicity is $\langle n_{ch} \rangle = 81$, whereas for the new models it is in the range 87–89.

the CM energy (at sufficiently high energies that energy–momentum conservation effects can be neglected). In the new model, the amount of primordial k_{\perp} given to initiators depends on the Q^2 of their associated hard scattering. With the increased phase space at the LHC, more primordial k_{\perp} is thus imparted by recoil effects to the junction baryon than at the Tevatron, hence the p_{\perp} spectrum becomes harder.

Also the junction baryon longitudinal migration shows some difference. Comparing the Tevatron junction baryon rapidity distribution, Fig. 26b above, with the LHC one, Fig. 32, we may distinguish two components. One is the peak at large rapidities, which corresponds to an (effective) diquark fragmentation and which is only shifted outwards in rapidity relative to the Tevatron by the increased energy. The other is the tail to central rapidities, which corresponds to baryon stopping. This tail *does* increase with energy, following the increase in the average number of interactions.

Finally, we show the $\langle p_{\perp} \rangle(n_{ch})$ distributions in Fig. 33. The same qualitative behaviour as at the Tevatron is apparent: the new models exhibit an earlier rise to a lower plateau, as compared to Tune A. Again, it is premature to draw

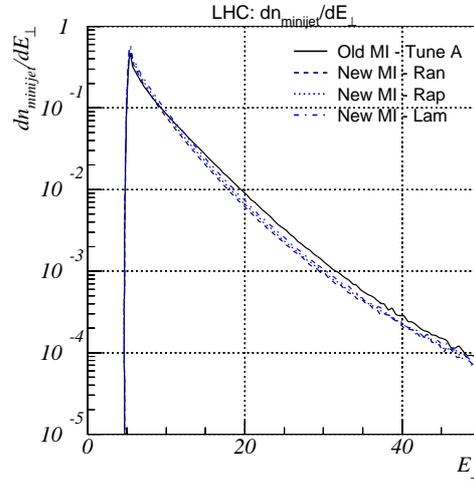


Figure 30: The number of jets as a function of jet E_{\perp} (for $E_{\perp} > 5\text{GeV}$) in minijet collisions at the LHC. Results are shown for Tune A (solid), Ran (dashed), Rap (dotted), and Lam (dash-dotted), as defined in Table 1.

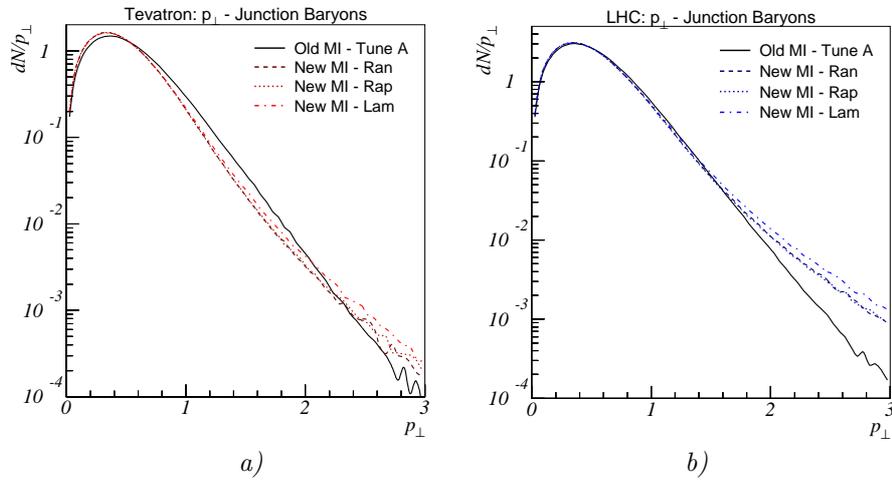


Figure 31: Junction baryon p_{\perp} spectrum at *a)* the Tevatron and *b)* the LHC. Results are shown for Tune A (solid), Ran (dashed), Rap (dotted), and Lam (dash-dotted), as defined in Table 1.

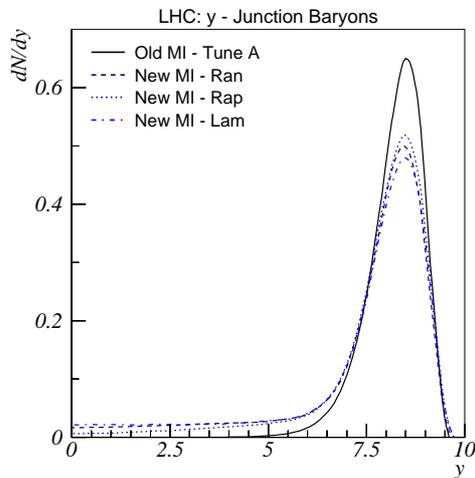


Figure 32: Junction baryon rapidity distributions at the LHC. Note: at the LHC *both* beam baryon numbers are included in the figure, whereas in the Tevatron plots, Fig. 26, the antibaryon number is not. Results are shown for Tune A (solid), Ran (dashed), Rap (dotted), and Lam (dash-dotted), as defined in Table 1.

any strong conclusions, in view of the still simple-minded description of the colour flow that we have included here. Further and more detailed studies of possible colour correlation mechanisms in hadronic collisions will be required in order to fully understand these aspects.

7 Conclusion and Outlook

Only in the last few years have multiple interactions gone from being a scientific curiosity, by most assumed relevant only for some rare topologies of four-jet events, to being accepted as the key element for understanding the structure of underlying events. However, this leaves a lot of questions to be addressed, such as:

- (i) What is the detailed mechanism and functional form of the dampening of the perturbative cross section at small p_{\perp} ?
- (ii) What is the energy dependence of the mechanism(s) involved?
- (iii) How is the internal structure of the proton reflected in an impact-parameter-dependent multiple interactions rate, as manifested e.g. in jet pedestal effects?
- (iv) How can the set of colliding partons from a hadron be described in terms

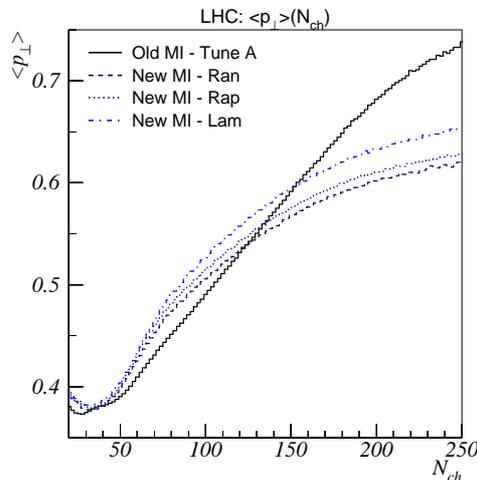


Figure 33: Average p_{\perp} as a function of charged multiplicity for min-bias collisions at the LHC. Results are shown for Tune A (solid), Ran (dashed), Rap (dotted), and Lam (dash-dotted), as defined in Table 1.

of correlated multiparton distribution functions of flavours and longitudinal momenta?

(v) How does a set of initial partons at some low perturbative cutoff scale evolve into such a set of colliding partons? Is standard DGLAP evolution sufficient, or must BFKL/CCFM effects be taken into account?

(vi) How would the set of initiators correlate with the flavour content of, and the longitudinal momentum sharing inside, the left-behind beam remnant?

(vii) How are the initiator and remnant partons correlated by confinement effects, e.g. in primordial k_{\perp} ?

(viii) How are all produced partons, both the interacting and the beam-remnant ones, correlated in colour? Is the large number-of-colours limit relevant, wherein partons can be hooked up into strings representing a linear confinement force?

(ix) How is the original baryon number of an incoming proton reflected in the colour topology?

(x) To what extent would a framework with independently fragmenting string systems, as defined from the colour topology, be modified by the space-time overlap of several strings?

Tentative answers to some of the questions are provided by the Tune A of the PYTHIA multiple interactions framework. Thus we now believe that:

- The matter overlap when two hadrons collide can be described by an impact-parameter dependence more spiked than a Gaussian but less so than an expo-

nential.

- The $p_{\perp 0}$ regularization scale does increase with energy.
- The colours of final-state partons are not random but correlated, somehow, to give a reduced string length.

This still leaves many questions unanswered. Worse, existing event generators would not even address many of the relevant issues, at least not in a deliberate or realistic fashion. In this article we have therefore tried to take the next step towards a better understanding of the structure of a hadronic event, addressing several of the points above. This in particular has concerned the correlations between initiator and remnant partons in the hadron beams, in terms of flavour, longitudinal and transverse momenta. Colour correlations have also been studied, and here it appears that the final-state partons need be involved as well. The complexity of the colour issues is tremendous, however, and we do not consider the studies finished in this area.

A specific new topic addressed is that of baryon number flow. Data from hadronic collisions, and even more from heavy-ion ones, show large excesses of baryon over antibaryon production in the central rapidity region of events [84], suggesting a significant influx of baryon number from the high-rapidity colliding beams, more than would be expected from standard quark/diquark fragmentation models. When the junction is introduced as a topological feature of the colour field in the baryon, however, the fate of the baryon number of an incoming beam particle may partly or wholly decouple from that of the valence quarks [1, 85]. We have here demonstrated that the junction topology in combination with multiple interactions can induce quite large rapidity shifts, of the desired kind.

The problem may actually be the opposite, i.e. not to move the baryon number by too much. To this end, we have assumed a suppression of interactions that affect several of the three colour chains that connect the valence quarks to the junction. This could be an impact-parameter-related effect, that not the whole proton is involved in the hard processes. If so, the suppression should be less pronounced in heavy-ion collisions, where the interactions of a proton with several nucleons in the other nucleus could occur at different positions in the transverse plane and thereby affect different chains. Obviously it would be a major undertaking to construct a complete model for heavy-ion collisions to study these ideas, but we hope in the future to be able to present a simple study of the baryon number flow.

Another open issue is that of intertwined initial-state showers, whereby two seemingly unrelated partons, each undergoing a hard scattering, reconstruct back to come from a common shower ancestor. With the new p_{\perp} -ordered showers now being implemented in PYTHIA [86] we intend to introduce enough flexibility that such issues could be addressed.

This will also further constrain the initial-state colour flow. The possibility of final-state colour reconnections remains, however, and has been pro-

posed as a mechanism to introduce diffractive topologies in a number of processes [87]. One here needs to better understand how much reconnections are allowed/required, and of what character.

We see that much work remains, before the physics of the underlying event is truly understood. Progress will not be possible without a constructive dialogue between theory and experiment. We have frequently had reason to mention Tune A as a role model here, because it offers a convenient reference that more sophisticated models can be tested against, without the need to know the details of the CDF detector. However, only a few distributions went into the tune, and so we do not know what to aim for in many other respects.

To give one specific example, it would be valuable to have information on the ‘lumpiness’ of the underlying event, such as n -jet rates as a function of some jet resolution parameter, similarly to e^+e^- -annihilation QCD analyses. One would there hope for an intermediate resolution region, between the coarse one that is dominated by the perturbative QCD structure and the fine one that mainly is sensitive to hadronization details, where the structure of the multiple interactions would play a key role. An understanding of this lumpiness is related to the fluctuations in the jet pedestal, and thereby to the smearing of jet energies in SUSY searches, say. It all hangs together . . .

In summary, striving for a better understanding of the physics of the underlying event is both interesting and useful. Interesting because it forces us to consider many issues normally swept under the carpet, and to confront dramatically different scenarios. Useful because it ties in with so many other physics analyses at hadron colliders. So there is plenty of interesting and useful work ahead of us before the picture has clarified completely!

Acknowledgments

The authors gratefully acknowledge the stimulating atmosphere at the Les Houches 2003 Workshop on Physics at TeV Colliders, the CERN Workshop on Monte Carlo tools for the LHC, and the Collider Physics 2004 Workshop at KITP, UCSB. We have benefitted from discussions with R.D. Field, J. Huston, A. Moraes, and many more. We are also grateful to the NorduGRID project, for use of computing resources.

This research was supported in part by the National Science Foundation under Grant No. PHY99-07949, and by The Royal Physiographic Society in Lund.

References

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

- [2] C. Seez, CMS TN/93–115;
CMS Collaboration, Technical Proposal, CERN/LHCC 94–38, p. 50,
ECAL Technical Design Report, CERN/LHCC 97–33, p. 330;
U. Egede, Ph.D. Thesis, Lund University LUNFD6/(NFFL–7150) 1997,
ISBN 91–628–2804–5;
ATLAS Collaboration, Detector and Physics Performance Technical De-
sign Report, Vol. I, CERN/LHCC 99–14, p. 228
- [3] T. Sjöstrand and P.Z. Skands, Nucl. Phys. **B659** (2003) 243
- [4] B. Andersson, G. Gustafson, G. Ingelman and T. Sjöstrand, Phys. Rep. **97** (1983) 31;
B. Andersson, ‘The Lund Model’ (Cambridge University Press, 1998)
- [5] P. Skands and T. Sjöstrand, in the proceedings of HEP 2003, Eur. Phys. J. C Direct, DOI: 10.1140/epjcd/s2003-03-520-7 [hep-ph/0310315];
T. Sjöstrand and P. Skands [hep-ph/0401060], in the proceedings of 3rd Les Houches Workshop on Physics at TeV Colliders, M. Dobbs et al. [hep-ph/0403100]
- [6] CTEQ Collaboration, H.L. Lai et al., Eur. Phys. J. **C12** (2000) 375
- [7] A. Donnachie and P.V. Landshoff, Phys. Lett. **B296** (1992) 227
- [8] J. Dischler and T. Sjöstrand, EPJdirect **C2** (2001) 1
- [9] T. Sjöstrand, P. Edén, C. Friberg, L. Lönnblad, G. Miu, S. Mrenna and E. Norrbin, Computer Phys. Commun. **135** (2001) 238;
T. Sjöstrand, L. Lönnblad, S. Mrenna and P. Skands, LU TP 03-38 [hep-ph/0308153]
- [10] G.A. Schuler and T. Sjöstrand, Phys. Rev. **D49** (1994) 2257
- [11] V.V. Sudakov, Zh.E.T.F. **30** (1956) 87 (Sov. Phys. J.E.T.P. **30** (1956) 65)
- [12] G.A. Schuler and T. Sjöstrand, Nucl. Phys. **B407** (1993) 539, Z. Phys. **C73** (1997) 677;
C. Friberg and T. Sjöstrand, JHEP **09** (2000) 10
- [13] H.-M. Chan et al., Nucl. Phys. **B86** (1975) 479, **B92** (1975) 13;
G.F. Chew and C. Rosenzweig, Phys. Rep. **41** (1978) 263
- [14] A. Capella, U. Sukhatme, C.-I. Tan and J. Tran Thanh Van, Phys. Lett. **81B** (1979) 68, Phys. Rep. **236** (1994) 225;
H. Minaka, Phys. Rev. **D20** (1979) 1656;
G. Cohen-Tannoudji et al., Phys. Rev. **D21** (1980) 2699;
A. Capella and J. Tran Thanh Van, Z. Phys. **C10** (1981) 249, Phys. Lett

- 114B** (1982) 450;
K. Fialkowski and A. Kotanski, Phys. Lett. **107B** (1981) 132, **115B** (1982) 425;
P. Aurenche and F.W. Bopp, Phys. Lett. **114B** (1982) 363;
A.B. Kaidalov, Phys. Lett. **116B** (1982) 459;
A.B. Kaidalov and K.A. Ter Martirosyan, Phys. Lett. **117B** (1982) 247;
P. Aurenche, F.W. Bopp and J. Ranft, Z. Phys. **C23** (1984) 67, **C26** (1984) 279, Phys. Rev. **D33** (1986) 1867;
V.M. Chudakov and V.V. Lugovoi, Z. Phys. **C59** (1993) 511
- [15] L.V. Gribov, E.M. Levin and M.G. Ryskin, Phys. Rept. **100** (1983) 1;
E.M. Levin and M.G. Ryskin, Phys. Rep. **189** (1990) 267
- [16] V.N. Gribov, Sov. Phys. JETP **26** (1968) 414;
V.A. Abramovski, O.V. Kancheli and V.N. Gribov, Soviet J. Nucl. Phys. **18** (1974) 308;
G. Veneziano, Nucl. Phys. **B74** (1974) 365, **B117** (1976) 519;
J. Bartels and M.G. Ryskin, Z. Phys. **C76** (1997) 241
- [17] H. Baer, F.E. Paige, S.D. Protopopescu, X. Tata, hep-ph/0001086
- [18] P. Aurenche, F.W. Bopp, A. Capella, J. Kwiecinski, M. Maire, J. Ranft and J. Tran Thanh Van, Phys. Rev. **D45** (1992) 92;
P. Aurenche, F.W. Bopp, R. Engel, D. Pertermann, J. Ranft and S. Roesler, Computer Phys. Commun. **83** (1994) 107
- [19] R. Engel, Z. Phys. **C66** (1995) 203;
R. Engel, J. Ranft, Phys. Rev. **D54** (1996) 4244
- [20] J. Ranft, Phys. Rev. **D51** (1995) 64;
S. Roesler, R. Engel and J. Ranft, hep-ph/0012252
- [21] P.V. Landshoff and J.C. Polkinghorne, Phys. Rev. **D18** (1978) 3344;
C. Goebel, F. Halzen and D.M. Scott, Phys. Rev. **D22** (1980) 2789;
M. Mangano, Z. Phys. **C42** (1989) 331
- [22] F. Takagi, Phys. Rev. Lett. **43** (1979) 1296;
N. Paver and D. Treleani, Nuovo Cimento **70A** (1982) 215, Phys. Lett. **146B** (1984) 252, Z. Phys. **C28** (1985) 187;
B. Humpert, Phys. Lett. **131B** (1983) 461;
B. Humpert and R. Odorico, Phys. Lett. **154B** (1985) 211;
G. Pancheri and Y.N. Srivastava, Phys. Lett. **B182** (1986) 199;
N. Brown, Mod. Phys. Lett. **A4** (1989) 2447
- [23] R.E. Ecclestone and D.M. Scott, Z. Phys. **C19** (1983) 29;
B. Humpert, Phys. Lett. **135B** (1984) 179;

- M. Mekhfi, Phys. Rev. **D32** (1985) 2371;
F. Halzen, P. Hoyer and W.J. Stirling, Phys. Lett. **B188** (1987) 375;
R.W. Robinett, Phys. Lett. **B230** (1989) 153;
R.M. Godbole, S. Gupta and J. Lindfors, Z. Phys. **C47** (1990) 69;
M. Drees and T. Han, Phys. Rev. Lett. **77** (1996) 4142;
A. Del Fabbro and D. Treleani, Phys. Rev. **D61** (2000) 077502;
A. Kulesza and W.J. Stirling, Phys. Lett. **B475** (2000) 168
- [24] T.T. Chou and C.N. Yang, Phys. Rev. **170** (1968) 1591, **D19** (1979) 3268;
C. Bourrely, J. Soffer and T.T. Wu, Nucl.Phys. **B247** (1984) 15; Eur. Phys. J. **C28** (2003) 97
P. L'Heureux, B. Margolis and P. Valin, Phys. Rev. **D32** (1985) 1681
- [25] L. Durand and P. Hong, Phys. Rev. Lett. **58** (1987) 303;
A. Capella, J. Kwiecinski and J. Tran Thanh Van, Phys. Rev. Lett. **58** (1987) 2015;
M.M. Block, F. Halzen and B. Margolis, Phys. Rev. **D45** (1992) 839;
R.M. Godbole and G. Pancheri, Eur. Phys. J. **C19** (2001) 129
- [26] J. Kuti and V.F. Weisskopf, Phys. Rev. **D4** (1971) 3418;
K. Konishi, A. Ukawa and G. Veneziano, Nucl. Phys. **B157** (1979) 45;
R. Kirschner, Phys. Lett. **84B** (1979) 266;
H.R. Gerhold, Nuovo Cimento **59A** (1980) 373;
V.P. Shelest, A.M. Snigirev and G.M. Zinovjev, Phys. Lett. **113B** (1982) 325;
A.M. Snigirev, Phys. Rev. **D68** (2003) 114012
- [27] M. Mekhfi, Phys. Rev. **D32** (1985) 2380;
M. Mekhfi and X. Artru, Phys. Rev. **D37** (1988) 2618
- [28] A. Bialas and E. Bialas, Acta. Phys. Pol. **B5** (1974) 373;
T.T. Chou and C.N. Yang, Phys. Rev. **D32** (1985) 1692;
S. Barshay, Z. Phys. **C32** (1986) 513;
W.R. Chen and R.C. Hwa, Phys. Rev. **D36** (1987) 760
- [29] A.H. Mueller and J. Qiu, Nucl. Phys. **B268** (1986) 427
- [30] G. Corcella, I.G. Knowles, G. Marchesini, S. Moretti, K. Odagiri, P. Richardson, M.H. Seymour and B.R. Webber, JHEP **0101** (2001) 010, hep-ph/0210213
- [31] UA5 Collaboration, G.J. Alner et al., Nucl. Phys. **B291** (1987) 445
- [32] J.M. Butterworth, J.R. Forshaw and M.H. Seymour, Z. Phys. **C72** (1996) 637
- [33] I. Borozan and M.H. Seymour, JHEP **0209** (2002) 015

- [34] E.A. Kuraev, L.N. Lipatov and V.S. Fadin, *Sov. Phys. JETP* **45** (1977) 199;
I.I. Balitsky and L.N. Lipatov, *Sov. J. Nucl. Phys.* **28** (1978) 822
- [35] M. Ciafaloni, *Nucl. Phys.* **B296** (1988) 49;
S. Catani, F. Fiorani and G. Marchesini, *Nucl. Phys.* **B336** (1990) 18
- [36] B. Andersson, G. Gustafson and J. Samuelsson, *Nucl. Phys.* **B467** (1996) 443;
B. Andersson, G. Gustafson and H. Kharraziha, *Phys. Rev.* **D57** (1998) 5543;
H. Kharraziha and L. Lönnblad, *JHEP* **03** (1998) 006
- [37] G. Gustafson, L. Lönnblad and G. Miu, *JHEP* **09** (2002) 005
- [38] G. Gustafson, L. Lönnblad and G. Miu, *Phys. Rev.* **D67** (2003) 034020
- [39] K. Kajantie, P.V. Landshoff and J. Lindfors, *Phys. Rev. Lett.* **59** (1987) 2527;
K.J. Eskola, K. Kajantie and J. Lindfors, *Nucl. Phys.* **B323** (1989) 37;
K.J. Eskola, K. Kajantie and K. Tuominen, *Nucl. Phys.* **A700** (2002) 509
- [40] AFS Collaboration, T. Åkesson et al., *Z. Phys.* **C34** (1987) 163
- [41] UA2 Collaboration, J. Alitti et al., *Phys. Lett.* **B268** (1991) 145
- [42] CDF Collaboration, F. Abe et al., *Phys. Rev.* **D47** (1993) 4857
- [43] CDF Collaboration, F. Abe et al., *Phys. Rev. Lett.* **79** (1997) 584, *Phys. Rev.* **D56** (1997) 3811
- [44] ZEUS Collaboration, C. Gwenlan, *Acta Phys. Polon.* **B33** (2002) 3123
- [45] D0 Collaboration, V.M. Abazov et al., *Phys. Rev.* **D67** (2003) 052001
- [46] UA1 Collaboration, C.-E. Wulz, in proceedings of the 22nd Rencontres de Moriond, Les Arcs, France, 15-21 March 1987
- [47] Z. Koba, H.B. Nielsen and P. Olesen, *Nucl. Phys.* **B40** (1972) 317
- [48] UA5 Collaboration, G.J. Alner et al., *Phys. Lett.* **138B** (1984) 304;
UA5 Collaboration, R.E. Ansorge et al., *Z. Phys.* **C43** (1989) 357
- [49] E735 Collaboration, T. Alexopoulos et al., *Phys. Lett.* **B435** (1998) 453
- [50] E.D. Malaza and B.R. Webber, *Nucl. Phys.* **B267** (1986) 702;
E.D. Malaza, *Z. Phys.* **C31** (1986) 143

- [51] ZEUS Collaboration, M. Derrick et al., *Z. Phys.* **C68** (1995) 29;
H1 Collaboration, S. Aid et al., contribution 479 to EPS HEP95, Brussels;
T. Sjöstrand, *J. Phys.* **G22** (1996) 709;
H1 Collaboration, Abstract 1102, submitted to ICHEP02, Amsterdam
- [52] R608 Collaboration, A.M. Smith et al., *Phys. Lett.* **B163** (1985) 267
- [53] UA5 Collaboration, R.E. Ansorge et al., *Z. Phys.* **C37** (1988) 191;
E735 Collaboration, T. Alexopoulos et al., *Phys. Lett.* **B353** (1995) 155
- [54] UA1 Collaboration, C. Albajar et al., *Nucl. Phys.* **B309** (1988) 405
- [55] CDF Collaboration, T. Affolder et al., *Phys. Rev.* **D65** (2002) 092002
- [56] R.D. Field (CDF Collaboration), in the proceedings of Snowmass 2001, eConf C010630 (2001) P501 [hep-ph/0201192], CDF Note 6403, talks available from webpage <http://www.phys.ufl.edu/~rfield/cdf/>;
V. Tano (CDF Collaboration), to appear in the proceedings of the 37th Rencontres de Moriond, Les Arcs, France, 16–23 March 2002 [hep-ex/0205023]
- [57] H1 Collaboration, S. Aid et al., *Z. Physik* **C70** (1996) 17
- [58] A. Breakstone et al., *Phys. Lett.* **132B** (1983) 463;
UA1 Collaboration, C. Albajar et al., *Nucl. Phys.* **B335** (1990) 261;
E735 Collaboration, T. Alexopoulos et al., *Phys. Lett.* **B336** (1994) 599;
CDF Collaboration, D. Acosta et al., *Phys. Rev.* **D65** (2002) 072005
- [59] UA1 Collaboration, *Phys. Lett.* **132B** (1983) 214;
D0 Collaboration, S. Abachi et al., *Phys. Lett.* **B357** (1995) 500;
ZEUS Collaboration, J. Breitweg et al., *Eur. Phys. J.* **C1** (1998) 109
- [60] CDF Collaboration, F. Abe et al., *Phys. Rev.* **D59** (1999) 032001;
P. Bartalini et al., in the proceedings of the Workshop on Standard Model physics (and more) at the LHC, CERN 2000-004, p.293
- [61] R.D. Field, presentations at the ‘Matrix Element and Monte Carlo Tuning Workshop’, Fermilab, 4 October 2002 and 29–30 April 2003, talks available from webpage <http://cepa.fnal.gov/CPD/MCTuning/>, and further recent talks available from <http://www.phys.ufl.edu/~rfield/cdf/>
- [62] A. Kharchilava, presentation at the ‘Matrix Element and Monte Carlo Tuning Workshop’, Fermilab, 4 October 2002, talk available from webpage <http://cepa.fnal.gov/CPD/MCTuning/>
- [63] A. Moraes, C. Buttar, I. Dawson and P. Hodgson, ATLAS internal notes; notes and talks available from webpage <http://amoraes.home.cern.ch/amoraes/>

- [64] J. M. Butterworth and S. Butterworth, *Comput. Phys. Commun.* **153** (2003) 164;
see also <http://jetweb.hep.ucl.ac.uk>
- [65] P.V. Landshoff, *Nucl. Phys. Proc. Suppl.* **99A** (2001) 311
- [66] Yu.L. Dokshitzer, V.A.Khoze, A.H. Mueller and S.I. Troyan, ‘Basics of Perturbative QCD’ (Editions Frontières, 1991)
- [67] T. Sjöstrand, *Phys. Lett.* **157B** (1985) 321;
M. Bengtsson, T. Sjöstrand and M. van Zijl, *Z. Phys.* **C32** (1986) 67
- [68] C. Balázs, J. Huston and I. Puljak, *Phys. Rev.* **D63** (2001) 014021
- [69] E. Thomé, master’s thesis, Lund University, LU TP 04–01 [hep-ph/0401121];
J. Huston, I. Puljak, T. Sjöstrand, E. Thomé, LU TP 04–07 [hep-ph/0401145]
- [70] EMC Collaboration, M. Arneodo et al., *Z. Physik* **C36** (1987) 527;
L. Apanasevich et al., *Phys. Rev.* **D59** (1999) 074007
- [71] G. ’t Hooft, *Nucl. Phys.* **B72** (1974) 461
- [72] B. Andersson, G. Gustafson and T. Sjöstrand, *Phys. Lett.* **B94** (1980) 211;
JADE Collaboration, W. Bartel et al., *Phys. Lett.* **101B** (1981) 129, *Z. Phys.* **C21** (1983) 37, *Phys. Lett.* **134B** (1984) 275;
TPC/2 γ Collaboration, H. Aihara et al., *Z. Phys.* **C28** (1985) 31;
TASSO Collaboration, M. Althoff et al., *Z. Phys.* **C29** (1985) 29
- [73] L. Montanet, G.C. Rossi and G. Veneziano, *Phys. Rep.* **63** (1980) 149
- [74] X. Artru, *Nucl. Phys.* **B85** (1975) 442
- [75] M. Imachi, S. Otsuki and F. Toyoda, *Progr. Theor. Phys.* **54** (1975) 280;
Y. Igarashi et al., *Progr. Theor. Phys. Suppl.* **63** (1978) 49;
G.C. Rossi and G. Veneziano, *Nucl. Phys.* **B123** (1977) 507
- [76] G. ’t Hooft, *Physica Scripta* **25** (1982) 133
- [77] P. Skands, in *Proceedings of the 10th International Conference on Supersymmetry and Unification of Fundamental Interactions*, DESY, Hamburg, 17–23 Jun 2002, eds. P. Nath, P.M. Zerwas and C. Grosche (Hamburg, DESY, 2002), vol. 2, p. 831 [hep-ph/0209199]
- [78] H.-U. Bengtsson, *Computer Physics Commun.* **31** (1984) 323

-
- [79] N. van Remortel, hep-ex/0305014, to appear in the proceedings of 38th Rencontres de Moriond on QCD and High-Energy Hadronic Interactions, Les Arcs, Savoie, France, 22-29 Mar 2003;
A. Straessner, hep-ex/0304031, *ibid.*
- [80] UA1 Collaboration, C. Albajar et al., Phys. Lett. **B226** (1989) 410;
E735 Collaboration, T. Alexopoulos et al., Phys. Rev. **D48** (1993) 1931
- [81] B. Andersson, G. Gustafson and B. Söderberg, Z. Phys. **C20** (1983) 317;
B. Andersson, S. Mohanty and F. Söderberg, Nucl. Phys. **B646** (2002) 102
- [82] E. Boos et al., in the proceedings of the Workshop on Physics at TeV Colliders, 21 May – 1 June 2001, Les Houches, France, eds. P. Aurenche et al. [hep-ph/0109068]
- [83] WA97 Collaboration, F. Antinori et al., Nucl. Phys. **A661** (1999) 130;
STAR Collaboration, C. Suires et al., Nucl. Phys. **A715** (2003) 470;
NA57 Collaboration, F. Antinori et al., J. Phys. **G30** (2004) S199
- [84] G. Belletini et al., Nuovo Cimento **A42** (1977) 85;
B. Alper et al., Nucl. Phys. **B100** (1975) 237;
A. Breakstone et al., Z. Physik **C28** (1985) 335;
H1 Collaboration, paper 556, submitted to ICHEP98, Vancouver, July 1998, and in preparation;
STAR Collaboration, C. Adler et al., Phys. Rev. Lett. **86** (2001) 4778;
PHENIX Collaboration, K. Adcox et al., Phys. Rev. Lett. **89** (2002) 092302;
PHOBOS Collaboration, B.B. Back et al., Phys. Rev. **C67** (2003) 021901;
BRAHMS Collaboration, I.G. Bearden et al., Phys. Rev. Lett. **90** (2003) 102301
- [85] D. Kharzeev, Phys. Lett. **B378** (1996) 238;
B. Kopeliovich and B. Povh, Z. Physik **C75** (1997) 693;
S.E. Vance, M. Gyulassy and X.-N. Wang, Phys. Lett. **B443** (1998) 45;
G.T. Garvey, B.Z. Kopeliovich, B. Povh, Comments Mod. Phys. **A2** (2001) 47
- [86] T. Sjöstrand, LU TP 04-05, submitted to the proceedings of the Workshop on Physics at TeV Colliders, Les Houches, France, 26 May – 6 June 2003 [hep-ph/0401061]
- [87] R. Enberg, G. Ingelman and N. Timneanu, Phys. Rev. **D64** (2001) 114015

**Transverse-Momentum-Ordered
Showers and Interleaved
Multiple Interactions**

Paper VI

Transverse-Momentum-Ordered Showers and Interleaved Multiple Interactions

T. Sjöstrand¹ and P. Z. Skands²

*Department of Theoretical Physics,
Lund University,
Sölvegatan 14A,
SE-223 62 Lund, Sweden*

Abstract

We propose a sophisticated framework for high-energy hadronic collisions, wherein different QCD physics processes are interleaved in a common sequence of falling transverse-momentum values. Thereby phase-space competition is introduced between multiple parton–parton interactions and initial-state radiation. As a first step we develop new transverse-momentum-ordered showers for initial- and final-state radiation, which should be of use also beyond the scope of the current article. These showers are then applied in the context of multiple interactions, and a few tests of the new model are presented. The article concludes with an outlook on further aspects, such as the possibility of a shower branching giving partons participating in two different interactions.

¹torbjorn@thep.lu.se

²peter.skands@thep.lu.se

1 Introduction

High-energy hadronic collisions offer a busy environment. The incoming hadrons seethe with activity as partons continuously branch and recombine. At the moment of collision, several partons from the two incoming hadrons may undergo interactions, that scatter the partons in different directions. The scattered partons may radiate, and all outgoing partons, including the beam remnants, hadronize in a correlated fashion to produce the observable high-multiplicity events. The physics involves a subtle blend of many perturbative and nonperturbative phenomena. No wonder that there is no simple, standard description to be offered!

What often saves the day is that most of the above activity is soft, i.e. confined to small transverse momenta p_{\perp} . When the processes of interest occur at large momentum transfers they therefore stand out, by producing jets, leptons or photons at large p_{\perp} . To first approximation, the rest of the activity, which we refer to as the *underlying event*, may then be disregarded. For precision studies, however, the problem remains: minijets from the underlying event may e.g. affect the jet energy calibration and the lepton and photon isolation criteria. Quite apart from the interesting challenge of better understanding the complex (semi-)soft processes for their own sake, this motivates an effort to investigate and model as well as possible the underlying event physics (when a selective trigger is used) and minimum-bias physics (for the inclusive sample of multihadronic events).

The basic building blocks needed to describe hadron-hadron collisions include hard-scattering matrix elements, parton density functions, initial- and final-state parton showers, and a hadronization scheme. Each of these deserve study in its own right, but additionally there is the question of how they should be combined. It is this latter aspect that we take aim at here. More specifically, we concentrate on the non-trivial interplay between multiple parton-parton interactions and initial-state parton showers, extending previous models for multiple interactions and developing new models for p_{\perp} -ordered initial- and final-state parton showers in the process.

A good starting point for the discussion is offered by Fig. 1. Based on the composite nature of hadrons we have here depicted *multiple interactions* (MI) between several pairs of incoming partons, see ref. [1] for a minireview. The structure of an incoming hadron is illustrated, with the p_{\perp} evolution of some partons from a nonperturbative border at $p_{\perp\text{min}}$ up to the different perturbative interactions. The $p_{\perp}^2 = \hat{t}\hat{u}/\hat{s}$ scale is a convenient measure of hardness, since the t (and u) channel gluon exchange processes $qq' \rightarrow qq'$, $qg \rightarrow qg$ and $gg \rightarrow gg$ dominate the cross section. One has to imagine a corresponding picture for the other hadron — omitted for clarity — with the two incoming sides joined at the interactions.

The next immediate issue that arises is how to describe hadronic objects

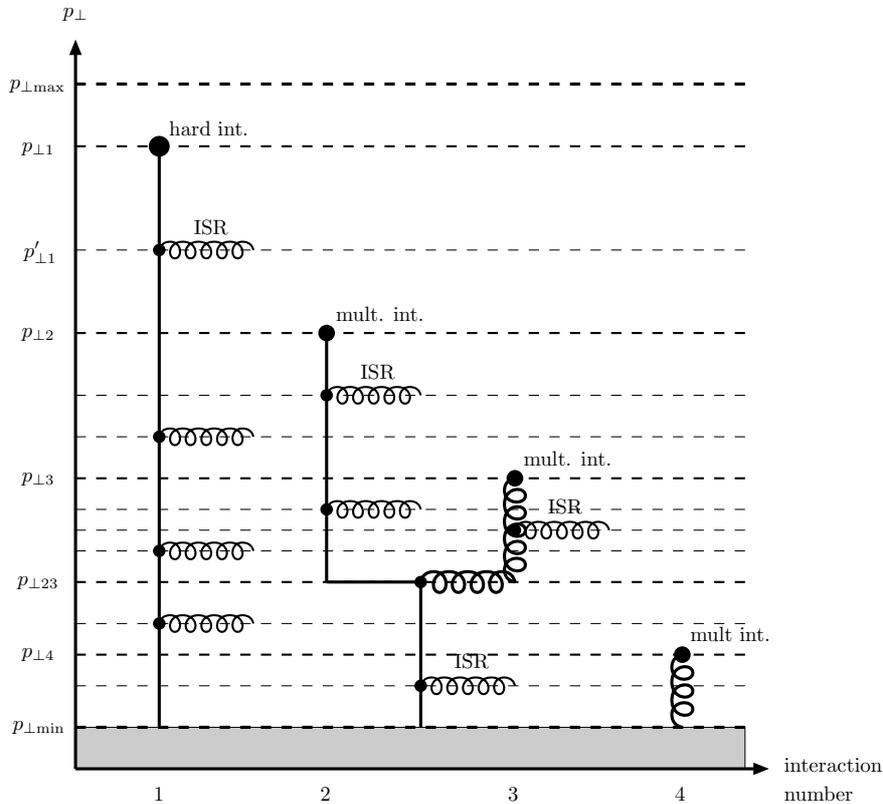


Figure 1: Schematic figure illustrating one incoming hadron in an event with a hard interaction occurring at $p_{\perp,1}$ and three further interactions at successively lower p_{\perp} scales, each associated with (the potentiality of) initial-state radiation, and further with the possibility of two interacting partons (2 and 3 here) having a common ancestor in the parton showers. Full lines represent quarks and spirals gluons. The vertical p_{\perp} scale is chosen for clarity rather than realism; most of the activity is concentrated to small p_{\perp} values.

under such conditions. In general, cross section calculations rely on parton density functions to describe the initial state. For the joint cross section of several simultaneous interactions one thus needs multi-parton densities, categorized by flavour content and fully differential in all x and $Q^2 \approx p_{\perp}^2$ values. Obviously such densities are almost entirely unconstrained, with neither data nor first-principles theory giving more than the roughest guidelines. To develop a realistic approximate framework, it is natural to consider first the hardest

interaction, which after all should be the most important one in terms of experimental consequences. Moreover, self-consistency ensures that this is also the interaction for which the standard ‘one-parton-inclusive’ pdf’s should be applicable; when averaging over all configurations of softer partons, the standard QCD phenomenology should be obtained for the ones participating in the hardest interaction, this being the way the standard parton densities have been measured. Thus it makes sense to order and study the interactions in a sequence of falling ‘hardness’, for which we shall here take p_{\perp} as our measure, i.e. we consider the interactions in a sequence $p_{\perp 1} > p_{\perp 2} > p_{\perp 3} > p_{\perp 4}$. The normal parton densities can then be used for the scattering at $p_{\perp 1}$, and correlation effects, known or estimated, can be introduced in the choice of ‘subsequent’ lower- p_{\perp} scatterings.

In ref. [1] we developed a new and sophisticated model to take into account such correlations in momentum and flavour. In particular, contrary to the earlier model described in ref. [2], the new model allows for more than one valence quark to be kicked out, and also takes into account the fact that sea quarks come in pairs. The beam remnant structure and colour flow topologies can become quite complicated, and so-called string junctions have to be handled, see [3].

In addition, the more sophisticated machinery allowed a more complete treatment of *initial-state radiation* (ISR) and *final-state radiation* (FSR). That is, each simple $2 \rightarrow 2$ interaction could be embedded in the center of a more complicated $2 \rightarrow n$ process, $n \geq 2$, where additional partons are produced by ISR or FSR. In order to avoid doublecounting, this additional radiation should be softer than the core $2 \rightarrow 2$ interaction. Here p_{\perp}^2 is again a convenient measure for hardness ordering, but not a unique one.

In this article, we introduce an additional interplay, between multiple parton interactions and ISR. ISR is the mechanism whereby parton densities evolve and become scale-dependent. The paradigm is that parton densities at a scale Q^2 , in our case identified with p_{\perp}^2 , probe the resolved partonic content at that scale. Therefore the issue of multi-parton densities is mixed in with the handling of ISR. For instance, if an ISR branching related to the first interaction occurs at a $p'_{\perp 1} < p_{\perp 1}$ then that reduces the available phase space for a second interaction at $p_{\perp 2} < p'_{\perp 1}$. In the complementary region $p_{\perp 2} > p'_{\perp 1}$, it is instead the momentum carried away by the second interaction that reduces the phase space for the ISR branching of the first. Thus, a consistent choice is to consider ISR (on both of the two incoming hadron sides) and MI in parallel, in one common sequence of decreasing p_{\perp} values, where the partonic structure at one p_{\perp} scale defines what is allowed at lower scales. Again this approach of *interleaved evolution* is intended to accurately reproduce measurements at p_{\perp} values corresponding to the hardest scales in the event, and fits well with the backwards evolution approach to ISR [4]. (One could have devised alternative procedures with forward evolution from lower to higher p_{\perp} values, which would

have offered a more intuitive physics picture, but with problems of its own.)

To the best of our knowledge, a scenario of this kind has never before been studied. In the early multiple interactions modelling [2] ISR and FSR was only included for the hardest interaction, and this before additional interactions were at all considered. In our more recent study [1] all interactions included ISR and FSR, but again separately for each interaction.

An additional difference is that, in our previous studies, spacelike (for ISR) or timelike (for FSR) virtuality was used as evolution and ordering variable in the showers. In the framework we shall present here, an essential ingredient is the use of p_{\perp} -ordered showers, such that the proposed competition between MI and ISR can be introduced in terms of a common ordering variable. We have therefore completed the rewriting begun in [5] of the existing PYTHIA showering algorithms [6] to p_{\perp} -ordering. These new models have interesting features in their own right, quite apart from the application to interleaved multiple interactions.

This article should be viewed as one step on the way towards a better understanding of hadronic physics, but not as the final word. Further issues abound. The downwards evolution in p_{\perp} may also reveal that two seemingly separately interacting partons actually have a common origin in the branching of a single parton at a lower p_{\perp} scale ($p_{\perp 23}$ in Fig. 1), and a single parton may scatter twice against partons in the other hadron. We shall refer to such possibilities specifically as *intertwined multiple interactions*, to distinguish them somewhat from the *interleaved* evolution that will be our main focus here.

In this article we begin, in Section 2, with a description of the new showering framework. This is followed, in Section 3, by a discussion on the model for interleaving MI and ISR, and a few results are presented in Section 4. The outlook in Section 5 contains a first estimate of the significance of the backward evolution joining several interactions. Finally Section 6 gives our conclusions.

2 New Transverse-Momentum-Ordered Showers

In this section we describe the new framework for timelike FSR and spacelike ISR in the context of a single hard-scattering process. We start by a brief review of the main existing showering algorithms, to introduce the basic terminology and ideas we will make use of. Thereafter the philosophy underlying the new algorithms is outlined. The more technical details are then described separately, first for timelike showers and then for spacelike ones, the latter as a rule being the more complicated.

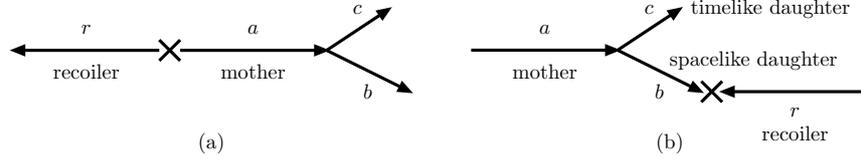


Figure 2: Schematic figure with our standard terminology for (a) a final-state and (b) an initial-state branching $a \rightarrow bc$, with a cross marking the central hard process and a recoiling parton r moving out to or coming in from the other side.

2.1 Shower minireview

In the shower approach, the evolution of a complex multi-parton final state is viewed as a succession of simple parton branchings. Thus a $2 \rightarrow n$ process can be viewed as consisting of a simple high-virtuality process, often $2 \rightarrow 2$, that approximately defines the directions and energies of the hardest jets of the process, combined with shower branchings at lower virtuality scales. The shower branchings thus add details to the simple answer, both by the production of additional jets and by a broadening of the existing ones. We distinguish between initial-state showers, whereby the incoming partons to the hard process build up increasingly spacelike virtualities Q^2 , and final-state showers, where outgoing partons, including the non-colliding partons emitted from the initial state, may have timelike virtualities Q^2 that decrease in the cascade down to on-shell partons.

To first order, both cascade types are governed by the same DGLAP evolution equations [7]

$$d\mathcal{P}_a(z, Q^2) = \frac{dQ^2}{Q^2} \frac{\alpha_s}{2\pi} P_{a \rightarrow bc}(z) dz, \quad (1)$$

expressing the differential probability that a ‘mother’ parton a will branch to two ‘daughter’ partons b and c , at a virtuality scale Q^2 , and with parton b taking a fraction z of the a energy, and c a fraction $1 - z$, cf. Fig.2. The splitting kernels $P_{a \rightarrow bc}(z)$ are (for massless quarks)

$$P_{q \rightarrow qg}(z) = \frac{4}{3} \frac{1+z^2}{1-z}, \quad (2)$$

$$P_{g \rightarrow gg}(z) = 3 \frac{(1-z(1-z))^2}{z(1-z)}, \quad (3)$$

$$P_{g \rightarrow q\bar{q}}(z) = \frac{n_f}{2} (z^2 + (1-z)^2), \quad (4)$$

where n_f is the number of quark flavours kinematically allowed. The kernels can be viewed as the universal collinear limit of the behaviour of rel-

evant matrix-element expressions. In such a context it is natural to associate Q^2 with $|m^2|$, the virtuality of an intermediate off-shell parton, since a $1/m^2$ comes from the propagator of the virtual particle. This is a free choice, however: if $Q^2 = f(z)m^2$, then for any (nice) function $f(z)$ it holds that $dQ^2/Q^2 dz = dm^2/m^2 dz$. At this stage, several equivalent choices are therefore possible.

Note that eq. (1) formally corresponds to the emission of an infinite number of partons. However, very soft and collinear gluons will not be resolved in an infrared safe fragmentation framework such as the string one [8], so we are free to introduce some effective Q_0 cut-off scale, of the order of 1 GeV or Λ_{QCD} , below which perturbative emissions need not be considered (to first approximation).

The remaining total emission probability is still normally above unity, which is allowed for an inclusive rate since several emissions can occur. For an exclusive parton shower it is then convenient to introduce a ‘time’ ordering, i.e. to decide which of the allowed emissions occur ‘first’. This is encompassed in the Sudakov form factor [9], expressing the probability that no emissions occur between the initial maximum scale Q_{max}^2 and a given Q^2 , and within limits $z_{\text{min}} < z < z_{\text{max}}$ that depend on the kinematics and the Q_0 cutoff,

$$\mathcal{P}_a^{\text{no}}(Q_{\text{max}}^2, Q^2) = \exp \left(- \int_{Q^2}^{Q_{\text{max}}^2} \int_{z_{\text{min}}}^{z_{\text{max}}} d\mathcal{P}_a(z', Q'^2) \right), \quad (5)$$

so that the differential probability for the first branching to occur at a $Q^2 = Q_a^2$ is given by $d\mathcal{P}_a(z, Q_a^2) \mathcal{P}_a^{\text{no}}(Q_{\text{max}}^2, Q_a^2)$. Once the parton a has branched, it is now the daughters b and c that can branch in their turn, with their Q_{max}^2 given by Q_a^2 , and so on until the cutoff scale is reached. Thus the shower builds up.

Obviously, at this stage different Q^2 choices are no longer equivalent: since a will only branch once, those regions of phase space considered at a later stage will be suppressed by a Sudakov factor relative to those considered earlier.

For ISR, the most commonly adopted approach is that of backwards evolution [4], wherein branchings are reconstructed backwards in time/virtuality from the hard interaction to the shower initiators. The starting point is the DGLAP equation for the b density

$$df_b(x, Q^2) = \frac{dQ^2}{Q^2} \frac{\alpha_s}{2\pi} \int \frac{dx'}{x'} f_a(x', Q^2) P_{a \rightarrow bc} \left(\frac{x}{x'} \right). \quad (6)$$

This expresses that, during a small increase dQ^2 there is a probability for parton a with momentum fraction x' to become resolved into parton b at $x = zx'$ and another parton c at $x' - x = (1-z)x'$. Correspondingly, in backwards evolution, during a decrease dQ^2 a parton b may become ‘unresolved’ into parton a . The relative probability $d\mathcal{P}_b$ for this to happen is given by the ratio df_b/f_b , which

translates into

$$d\mathcal{P}_b(x, Q^2) = \left| \frac{dQ^2}{Q^2} \right| \frac{\alpha_s}{2\pi} \int dz \frac{x' f_a(x', Q^2)}{x f_b(x, Q^2)} P_{a \rightarrow bc}(z). \quad (7)$$

Again, ordering the evolution in Q^2 implies that this ‘naive probability’ should be multiplied by the probability $\mathcal{P}_b^{\text{no}}(x, Q_{\text{max}}^2, Q^2)$ for no emissions to occur at scales higher than Q^2 , obtained from $d\mathcal{P}_b$ by exponentiation like in eq. (5). As for the timelike showers, additional sophistication can be added by coherence constraints and matrix-element merging, but ISR remains less well understood than FSR [22].

2.2 Existing approaches

Of the three most commonly used final-state shower algorithms, PYTHIA uses m^2 as evolution variable [10, 11], while HERWIG uses an energy-weighted emission angle, $E^2(1 - \cos\theta) \sim m^2/(z(1-z))$ [12], and ARIADNE a squared transverse momentum, $\sim z(1-z)m^2$ [13–15]. Thus the three programs give priority to emissions with large invariant mass, large emission angle and large transverse momentum, respectively.

The HERWIG algorithm makes angular ordering a direct part of the evolution process, and thereby correctly (in an azimuthal-angle-averaged sense) takes into account coherence effects in the emission of soft gluons [16]. Branchings are not ordered in hardness: often the first emission is that of a soft gluon at wide angles. The algorithm does not populate the full phase space but leaves a ‘dead zone’ in the hard three-jet region, that has to be filled up separately [17]. The kinematics of a shower is only constructed at the very end, after all emissions have been considered.

The PYTHIA algorithm is chosen such that the shower variables closely match the standard three-jet phase space in $e^+e^- \rightarrow q\bar{q}g$, and such that the shower slightly overpopulates the hard three-jet region, so that a simple rejection step can be used to obtain a smooth merging of all relevant first-order gluon-emission matrix elements with the shower description [11]. The mass-ordering of emissions is one possible definition of hardness-ordering. The main limitation of the algorithm is that it does not automatically include coherence effects. Therefore angular ordering is imposed by an additional veto, but then cuts away a bit too much of the soft-gluon phase space [18]. The kinematics of a branching is not constructed until the daughters have been evolved in their turn, so that their virtualities are also known.

The ARIADNE algorithm differs from the above two in that it is formulated in terms of dipoles, consisting of parton pairs, rather than in terms of individual partons. The two partons that make up a dipole may then collectively emit a gluon, causing the dipole to split in two. Thus the basic process is that of one dipole branching into two dipoles, rather than of one parton branching into two

partons. Emissions are ordered in terms of a decreasing transverse momentum, which automatically includes coherence effects [13], and also is a good measure of hardness. Kinematics can be constructed, in a Lorentz invariant fashion, immediately after each branching, with individual partons kept on mass shell at each stage. This makes it easy to stop and restart the shower at some intermediate p_{\perp} scale. The implementation of an (L)CKKW-style matching of matrix elements with parton showers [19] is therefore simplified, and in particular Sudakov factors can be generated dynamically to take into account the full kinematics of the branching history. A disadvantage is that $g \rightarrow q\bar{q}$ branchings do not fit naturally into a dipole framework, since they cannot be viewed as one dipole branching into two.

In experimental tests, e.g. compared with LEP data [21], the three final-state algorithms all offer acceptable descriptions. If HERWIG tends to fare the worst, it could partly reflect differences in the hadronization descriptions, where the HERWIG cluster approach is more simplistic than the PYTHIA string one, also used by ARIADNE. Among the latter two, ARIADNE tends to do somewhat better.

The above three programs also can be used for initial-state showers. For HERWIG the evolution variable is again angular-defined, and for PYTHIA now $Q^2 = -m^2$. Both programs make use of backwards evolution, as described above.

By contrast, the ARIADNE approach defines radiating dipoles spanned between the remnants and the hard scattering [23], and thereby cannot easily be related to the standard DGLAP formalism. LDCMC is a more sophisticated approach [24], based on forward evolution and unintegrated parton densities, and equivalent to the CCFM equations [25].

2.3 The new approach

In this article we wish to modify/replace the existing PYTHIA shower routines so that emissions are ordered in p_{\perp}^2 rather than in $Q^2 = \pm m^2$, and also include some of the good points of the dipole approach within the shower formalism. Specifically we

- retain the shower language of one parton branching into two, such that $g \rightarrow q\bar{q}$ appears on equal footing with other branchings,
- make use of a simplified p_{\perp}^2 as evolution variable, picked such that the translation $p_{\perp}^2 \leftrightarrow \pm m^2$ is trivial, thereby preserving all the sophistication of the existing matrix-element-merging,
- construct a preliminary kinematics directly after each branching, with currently unevolved partons explicitly on mass shell,
- define a recoil partner, ‘recoiler’, for each branching parton, ‘radiator’, to keep the total energy and momentum of the radiator+recoiler ‘dipole’ preserved whenever a parton previously put on mass shell is assigned a virtuality,

and

- ensure that the algorithms can be stopped and restarted at any given intermediate p_\perp scale without any change of the final result, so that they can be used for interleaving showers and multiple interactions (and also for (L)CKKW-style matching, although this will not be made use of here).

Transverse momentum definitions

So far, we have used p_\perp to denote a general kind of ‘transverse momentum’, without specifying further the details of which momentum we are talking about and which direction it is transverse to. It is now our purpose to specify more closely which precise definition(s) we have in mind, and to give a comparison to some other commonly encountered p_\perp definitions.

To specify a p_\perp suitable for a branching $a \rightarrow bc$, consider lightcone kinematics, $p^\pm = E \pm p_z$, for which $p^+p^- = m_\perp^2 = m^2 + p_\perp^2$. For a moving along the $+z$ axis, with $p_b^+ = zp_a^+$ and $p_c^+ = (1-z)p_a^+$, p^- conservation then gives

$$m_a^2 = \frac{m_b^2 + p_\perp^2}{z} + \frac{m_c^2 + p_\perp^2}{1-z} \quad (8)$$

or equivalently

$$p_\perp^2 = z(1-z)m_a^2 - (1-z)m_b^2 - zm_c^2 = p_{\perp\text{LLC}}^2. \quad (9)$$

For a timelike branching $Q^2 = m_a^2$ and $m_b = m_c = 0$, so then $p_{\perp\text{LLC}}^2 = z(1-z)Q^2$. For a spacelike branching $Q^2 = -m_b^2$ and $m_a = m_c = 0$, so instead $p_{\perp\text{LLC}}^2 = (1-z)Q^2$. We use these relations to define abstract evolution variables $p_{\perp\text{Levol}}^2 = z(1-z)Q^2$ or $= (1-z)Q^2$, in which to order the sequence of shower emissions.

However, this is *not* the z definition we will use to construct the kinematics of the branchings. For this, we interpret z to give the energy sharing between the daughters, in the rest frame of the radiator+recoiler system, $E_b = zE_a$ and $E_c = (1-z)E_a$. The latter z interpretation gives nice Lorentz invariance properties — energies in this frame are easily related to invariant masses, $2E_i/m_{ijk} = 1 - m_{jk}^2/m_{ijk}^2$ for the ijk three-parton configuration after the radiation — but gives more cumbersome kinematics relations, specifically for p_\perp . This is the reason we use the lightcone relations to define the evolution variable while we use the energy definition of z to construct the actual kinematics of the branchings.

The deliberate choice of maintaining this dichotomy can be better understood by examining a few different p_\perp definitions in common use, in particular those in clustering algorithms. To this end consider first the situation depicted in Fig. 3a: With the two particles massless, so that $E_1 = |\mathbf{p}_1|$ and $E_2 = |\mathbf{p}_2|$, the momentum transverse to the vector sum $\mathbf{p}_1 + \mathbf{p}_2$, which would correspond

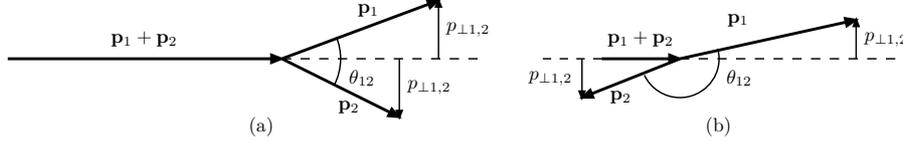


Figure 3: (a) Schematic figure of the clustering of two particles. (b) A topology with a large θ_{12} but a small $p_{\perp 1,2}$.

to the momentum of an imagined mother, is

$$p_{\perp} = \frac{|\mathbf{p}_1 \times \mathbf{p}_2|}{|\mathbf{p}_1 + \mathbf{p}_2|} = \frac{E_1 E_2 \sin \theta_{12}}{\sqrt{E_1^2 + E_2^2 + 2E_1 E_2 \cos \theta_{12}}} = p_{\perp 1,2} . \quad (10)$$

There is one troubling feature of this $p_{\perp 1,2}$: not only does it vanish when the opening angle θ_{12} goes to zero, but it also vanishes for $\theta_{12} \rightarrow \pi$ (unless $E_1 \equiv E_2$). Physically it is clear what is happening in this limit: the parton with larger energy is going along the $\mathbf{p}_1 + \mathbf{p}_2$ direction and the one with smaller energy is just opposite to it, Fig. 3b. In a clustering algorithm, where the idea is to combine ‘nearby’ particles, a measure with such a behaviour clearly is undesirable. Even when the starting point would be to have a p_{\perp} -related measure for small θ_{12} , we would prefer to have this measure increase monotonically for increasing θ_{12} , given fix E_1 and E_2 , and behave a bit more like the invariant mass at large angles. Therefore, in the LUCLUS algorithm [26], the replacements $\sin \theta_{12} \rightarrow 2 \sin(\theta_{12}/2)$ and $|\mathbf{p}_1 + \mathbf{p}_2| \rightarrow E_1 + E_2$ are performed, so that

$$p_{\perp} = \frac{|\mathbf{p}_1 \times \mathbf{p}_2|}{|\mathbf{p}_1 + \mathbf{p}_2|} \rightarrow \frac{E_1 E_2 2 \sin(\theta_{12}/2)}{E_1 + E_2} = p_{\perp L} . \quad (11)$$

But, since $\sin^2(\theta_{12}/2) = (1 - \cos \theta_{12})/2$, it also follows that

$$p_{\perp L}^2 = \frac{E_1}{E_1 + E_2} \frac{E_2}{E_1 + E_2} 2E_1 E_2 (1 - \cos \theta_{12}) \simeq z(1-z)m^2 = p_{\perp \text{evol}}^2 , \quad (12)$$

given our z definition in the shower as being one of energy sharing.

The $p_{\perp L}$ and $p_{\perp \text{evol}}$ are not completely equivalent: for the shower algorithm to be Lorentz invariant it is essential that the energies in the z definition are defined in the radiator+recoiler rest frame, whereas the LUCLUS algorithm normally would be applied in the rest frame of the event as a whole. Nevertheless, we gain some understanding why the choice of $p_{\perp \text{evol}}^2$ as evolution variable actually may be more physically meaningful than $p_{\perp 1,2}^2$. Specifically, for the emission of a gluon off a $q\bar{q}$ dipole, say, we retain the subdivision of radiation from the mass-ordered algorithm, roughly in proportions $1/m_{qg}^2 : 1/m_{\bar{q}g}^2$ for $q \rightarrow qg : \bar{q} \rightarrow \bar{q}g$. With the $p_{\perp 1,2}^2$ measure, q radiation close to the \bar{q} would not

be disfavoured, since also $\theta_{12} \rightarrow \pi$ would be classified as a collinear emission region.

The Durham clustering algorithm [27] is intended to represent the transverse momentum of the lower-energy parton relative to the direction of the higher-energy one, but again modified to give a sensible behaviour at large angles:

$$p_{\perp\text{rel}} = \min(E_1, E_2) \sin \theta_{12} \rightarrow \min(E_1, E_2) 2 \sin(\theta_{12}/2) = p_{\perp\text{D}} \quad (13)$$

Thereby it follows that

$$(p_{\perp\text{evol}} \simeq) p_{\perp\text{L}} = \frac{\max(E_1, E_2)}{E_1 + E_2} p_{\perp\text{D}} \quad (14)$$

so the two p_{\perp} measures never disagree by more than a factor of two, and coincide in the soft-gluon limit.

In the ARIADNE dipole emission approach, finally, the p_{\perp} is defined as the momentum of the emitted parton relative to the axis of the emitting partons [15]. For the emission of a soft parton 3 from the 1 and 2 recoiling parton dipole one can then derive

$$p_{\perp\text{A}}^2 = \frac{m_{13}^2 m_{23}^2}{m_{123}^2} . \quad (15)$$

When the $m^2 = m_{13}^2 \rightarrow 0$ limit is considered, this corresponds to $p_{\perp\text{A}}^2 \approx (1-z)m^2$, rather than the $p_{\perp\text{L}}^2 \approx z(1-z)m^2$. That is, for the soft-gluon limit $z \rightarrow 1$ the two measures agree, while they disagree in the hard-gluon limit $z \rightarrow 0$: $p_{\perp\text{A}}^2 \approx m^2 \gg zm^2 \approx p_{\perp\text{L}}^2$. It is not clear whether this difference by itself would have any visible consequences, but it illustrates that the meaning of ‘ p_{\perp} -ordered emission’ is not uniquely defined.

The new algorithms

Taking into account the above considerations, the basic strategy of the algorithms therefore can be summarized as follows:

1. Define the evolution variable $p_{\perp\text{evol}}^2$,

$$\text{FSR} : p_{\perp\text{evol}}^2 = z(1-z)Q^2 , \quad (16)$$

$$\text{ISR} : p_{\perp\text{evol}}^2 = (1-z)Q^2 . \quad (17)$$

2. Evolve all radiators downwards in $p_{\perp\text{evol}}^2$, from a $p_{\perp\text{max}}^2$ defined either by the hard process or by the preceding shower branching, to find trial branchings

according to the respective evolution equation,

$$\text{FSR : } d\mathcal{P}_a = \frac{dp_{\perp\text{evol}}^2}{p_{\perp\text{evol}}^2} \frac{\alpha_s(p_{\perp\text{evol}}^2)}{2\pi} P_{a\rightarrow bc}(z) dz \mathcal{P}_a^{\text{no}}(p_{\perp\text{max}}^2, p_{\perp\text{evol}}^2), \quad (18)$$

$$\text{ISR : } d\mathcal{P}_b = \frac{dp_{\perp\text{evol}}^2}{p_{\perp\text{evol}}^2} \frac{\alpha_s(p_{\perp\text{evol}}^2)}{2\pi} \frac{x' f_a(x', p_{\perp\text{evol}}^2)}{x f_b(x, p_{\perp\text{evol}}^2)} P_{a\rightarrow bc}(z) dz \times \mathcal{P}_b^{\text{no}}(x, p_{\perp\text{max}}^2, p_{\perp\text{evol}}^2). \quad (19)$$

Note that we have chosen $p_{\perp\text{evol}}^2$ as scale both for parton densities and α_s [28]. The Sudakov form factors are, as before, obtained by exponentiation of the respective real-emission expressions.

3. Select the radiator+recoiler set with the largest trial $p_{\perp\text{evol}}^2$ to undergo the next actual branching.
4. For this branching, use the picked $p_{\perp\text{evol}}^2$ and z values to derive the virtuality Q^2 ,

$$\text{FSR : } m_a^2 = Q^2 = \frac{p_{\perp\text{evol}}^2}{z(1-z)}, \quad (20)$$

$$\text{ISR : } -m_b^2 = Q^2 = \frac{p_{\perp\text{evol}}^2}{1-z}. \quad (21)$$

5. Construct kinematics based on Q^2 and z
 - a) in the radiator+recoiler rest frame,
 - b) defining z in terms of energy fractions, or equivalently mass ratios,
 - c) assuming that yet unbranched partons are on-shell and that the current two 'earliest' ISR partons are massless, and
 - d) shuffling energy-momentum from the recoiler as required.
6. Iterate towards lower $p_{\perp\text{evol}}^2$ until no further branchings are found above the lower cutoff scale $p_{\perp\text{min}}^2$.

We now proceed to fill in the details for the respective algorithms.

2.4 Timelike showers

The basic formalism

At each step of the evolution there is a set of partons that are candidates for further branching. Each such radiator defines dipoles together with one or several recoiler partons. Normally these recoilors are defined as the parton carrying the anticolour of the radiator, where colour indices in a cascade are traced in the $N_C \rightarrow \infty$ limit. A gluon, with both a colour and an anticolour index, thus has two partners, and the nominal emission rate is split evenly between these two. Since the kinematics constraints in the two radiator+recoiler dipoles normally will be different, the actual emission probabilities will not agree, however.

To illustrate, consider $e^+e^- \rightarrow \gamma^*/Z^0 \rightarrow q\bar{q}g$, where one gluon has already been radiated. The quark is then a radiator, with the gluon as recoiler, but also the gluon is a radiator with the quark as recoiler. Similarly for the antiquark–gluon pair. There is no colour dipole directly between the quark and the antiquark. On the other hand, we may also allow photon emission via the shower branching $q \rightarrow q\gamma$, similarly to $q \rightarrow qg$, and for such branchings indeed the quark and the antiquark are each other’s recoilers, while the uncharged gluon is not involved at all. In total, this configuration thus corresponds to six possible radiator+recoiler sets. Each of these are to be evolved downwards from the $p_{\perp\text{evol}}^2$ scale of the first gluon emission, and the one with largest new $p_{\perp\text{evol}}^2$ is chosen as the next evolution step to be realized. Thereafter the whole procedure is iterated, to produce one common sequence of branchings with $p_{\perp\text{max}} > p_{\perp 1} > p_{\perp 2} > \dots > p_{\perp\text{min}}$.

A special case is where a narrow coloured resonance is concerned, as for instance in top decay to bW^+ . Here, gluon emissions with energies above the width of the top should not change the top mass. They are constrained inside the top system. (In fact, when $b \rightarrow bg$, the other end of the colour dipole is rather defined by the decaying top itself.) Technically, the W^+ may then be chosen as the recoiler to the b , to ensure that the top mass remains unchanged. In this case all radiation is off the b , i.e. the system only contains one gluon radiator, and this is enough to reproduce the desired rate [11].

Once a recoiler has been assigned, the kinematics of a branching is suitably defined in the rest frame of the radiator+recoiler system, with the radiator a (recoiler r) rotated to move out along the $+z$ ($-z$) axis. Then one may define $m_{ar}^2 = (p_a + p_r)^2$. For massless partons, the introduction of an off-shell $Q^2 = m_a^2 = p_{\perp\text{evol}}^2/z(1-z)$ increases E_a from $m_{ar}/2$ to $(m_{ar}^2 + Q^2)/2m_{ar}$, with E_r reduced by the same amount. The two daughters share the energy according to $E_b = zE_a$ and $E_c = (1-z)E_a$. With the modified a still along the $+z$ axis, the transverse momentum of the two daughters then becomes

$$p_{\perp b,c}^2 = \frac{z(1-z)(m_{ar}^2 + Q^2)^2 - m_{ar}^2 Q^2}{(m_{ar}^2 - Q^2)^2} Q^2 \leq z(1-z)Q^2 = p_{\perp\text{evol}}^2. \quad (22)$$

The kinematics can now be completed, rotating and boosting the two daughters and the modified recoiler back to the original frame.

Note that $p_{\perp b,c}^2$ and $p_{\perp\text{evol}}^2$ always coincide for $z = 1/2$, and agree well over an increasing z range as $Q^2/m_{ar}^2 \rightarrow 0$. We have already explained why $p_{\perp\text{evol}}^2$ is a better evolution variable than $p_{\perp b,c}^2$. In addition, there are technical advantages: had evolution been performed in $p_{\perp b,c}^2$, the extraction of a Q^2 from $p_{\perp b,c}^2$ would require solving a third-degree equation, which would be messy and possibly give several solutions. The allowed z range would also be nontrivially defined. As it is now, the requirement $Q^2 < m_{ar}^2$ easily leads to a range

$z_{\min} < z < z_{\max}$ for $p_{\perp\text{evol}}^2$, with

$$z_{\min,\max} = \frac{1}{2} \left(1 \mp \sqrt{1 - \frac{p_{\perp\text{evol}}^2}{m_{ar}^2}} \right). \quad (23)$$

Once a trial $p_{\perp\text{evol}}^2$ and z has been picked, and thereby Q^2 is known, an acceptable solution has to be in the smaller range

$$z_{\min,\max} = \frac{1}{2} \left(1 \mp \frac{m_{ar}^2 - Q^2}{m_{ar}^2 + Q^2} \right) \quad (24)$$

for $p_{\perp b,c}^2 > 0$ to be valid.

It is the choice of a dipole-style phase space in conjunction with p_{\perp}^2 as evolution variable that ensures the angular ordering required for coherence [13,18].

Further details

(i) The colour topology of an event needs to be updated after each branching, so as to define possible recoilers for the next step of the evolution, and also for the subsequent hadronization. Most of this is trivial, since we work in the $N_C \rightarrow \infty$ limit: for $q \rightarrow qg$ the original quark colour is inherited by the gluon and a new colour dipole is created between the two daughters, while for $g \rightarrow q\bar{q}$ the (anti)quark takes the gluon (anti)colour. Somewhat more tricky is $g \rightarrow gg$, where two inequivalent possibilities exist. We here use the rewriting of the splitting kernel [14], $(1-z(1-z))^2/z(1-z) = (1+z^3)/(1-z) + (1+(1-z)^3)/z \simeq 2(1+z^3)/(1-z)$, to associate a $1-z$ picked according to the right-hand side with the energy fraction of the ‘radiated’ gluon that carries away the ‘radiating’ (anti)colour of the original gluon.

(ii) The above p_{\perp} equations have been written for the case of massless partons. It is straightforward to generalize to massive partons, however, starting from the formalism presented in ref. [11]. There it was shown that the natural variable for mass-ordered evolution of a parton a with on-shell mass $m_{a,0}$ is $Q^2 = m_a^2 - m_{a,0}^2$, since this reproduces relevant propagators. Now the generalization is

$$p_{\perp\text{evol}}^2 = z(1-z)(m_a^2 - m_{a,0}^2). \quad (25)$$

Furthermore, in the handling of kinematics, the z variable is reinterpreted to take into account masses [11].

(iii) Whether radiation off massive or massless partons is considered, matrix-element expressions are available for the one-gluon emission corrections in $a \rightarrow bc$ decays in the standard model and its minimal supersymmetric extension, say $\gamma^*/Z^0 \rightarrow q\bar{q}$ or $\tilde{g} \rightarrow \tilde{q}\bar{q}$ [11]. Since the shower overpopulates phase space relative to these expressions, a simple veto step can be used to smoothly merge

a matrix-element behaviour for hard non-collinear emissions with the shower picture for soft and collinear ones. When the $b + c$ system radiates repeatedly, the matrix-element corrections are applied to the system at the successively reduced energy. This ensures that a good account is given of the reduced radiation in the collinear region by mass effects. For $g \rightarrow q\bar{q}$ branchings, mass effects and subsequent gluon emissions off the quarks are given the same corrections as for $\gamma^* \rightarrow q\bar{q}$ branchings, i.e. disregarding the difference in colour structure.

(iv) Azimuthal φ angles are selected isotropically in $q \rightarrow qg$ branchings, but nonisotropically for $g \rightarrow gg$ and $g \rightarrow q\bar{q}$ to take into account gluon polarization effects [29]. Anisotropies from coherence conditions are not included explicitly, since some of that is implicitly generated by the dipole kinematics.

(v) We use a first-order $\alpha_s(p_\perp^2) = 12\pi/((33 - 2n_f) \ln(p_\perp^2/\Lambda_{(n_f)}^2))$, matched at the m_c and m_b mass thresholds, where default is $m_c = 1.5$ GeV and $m_b = 4.8$ GeV.

Algorithm tests

Ultimately, the usefulness of a shower algorithm is gauged by its ability to describe data. Obviously, we have checked that the results of the new routine qualitatively agree with the old program, which is known to describe data reasonably well. A more detailed study has been performed by G. Rudolph [30], who has compared our algorithm with ALEPH data at the Z^0 peak [21]. A tune to a set of event shapes and particle spectra gives a total χ^2 that is roughly 2/3 of the corresponding value for the old mass-ordered evolution, i.e. a marked improvement. Of the distributions considered, the only one that does not give a decent description is the single-particle $p_{\perp\text{out}}$ spectrum, i.e. the transverse momentum out of the event plane, in the region $p_{\perp\text{out}} > 0.7$ GeV. This is a common problem for showering algorithms, and in fact was even bigger in the mass-ordered one. With the exception of this region, the χ^2 per degree of freedom comes down to the order of unity, if one to the experimental statistical and systematical errors in quadrature adds an extra term of 1% of the value in each point. That is, it appears plausible that the overall quality of the algorithm is at the 1% level for most observables at the Z^0 peak.

Some of the tuned values have changed relative to the old algorithm. Specifically the first-order five-flavour Λ is roughly halved to 0.140 GeV, and the cutoff parameter is reduced from $m_{\text{min}} \approx 1.6$ GeV to $2p_{\perp\text{min}} \approx 0.6$ GeV. The former represents a real enough difference in the capability of the algorithms to populate the hard-emission region, while the latter is less easily interpreted and less crucial, since it deals with how best to match perturbative and nonperturbative physics, that is largely compensated by retuned hadronization parameters.

2.5 Spacelike showers

The basic formalism

At any resolution scale $p_{\perp\text{evol}}^2 = (1-z)Q^2$ the ISR algorithm will identify two initial partons, one from each incoming hadron, that are the mothers of the respective incoming cascade to the hard interaction. When the resolution scale is reduced, using backwards evolution according to eq. (19), either of these two partons may turn out to be the daughter b of a previous branching $a \rightarrow bc$. The (currently resolved) parton r on the other side of the event takes on the role of recoiler, needed for consistent reconstruction of the kinematics when the parton b previously considered massless now is assigned a spacelike virtuality $m_b^2 = -Q^2$. This redefinition should be performed in such a way that the invariant mass of the $b+r$ system is unchanged, since this mass corresponds to the set of outgoing partons already defined by the hard scattering and by partons emitted in previously considered branchings. The system will have to be rotated and boosted as a whole, however, to take into account that b not only acquires a virtuality but also a transverse momentum; if previously b was assumed to move along the event axis, now it is a that should do so.

At any step of the cascade, the massless mothers suitably should have four-momenta given by $p_i = x_i(\sqrt{s}/2)(1; 0, 0, \pm 1)$ in the rest frame of the two incoming beam particles, so that $\hat{s} = x_1 x_2 s$. If this relation is to be preserved in the $a \rightarrow bc$ branching, the $z = x_b/x_a$ should fulfil $z = m_{br}^2/m_{ar}^2 = (p_b + p_r)^2/(p_a + p_r)^2$. As we have already noted, z definitions in terms of squared mass ratios are easily related to energy sharing in the rest frame of the process. This is illustrated by explicit construction of the kinematics in the $a+r$ rest frame, assuming a moving along the $+z$ axis and c massless:

$$p_{a,r} = \frac{m_{ar}}{2} (1; 0, 0, \pm 1) , \quad (26)$$

$$p_b = \left(\frac{m_{ar}}{2} z; \sqrt{(1-z)Q^2 - \frac{Q^4}{m_{ar}^2}}, 0, \frac{m_{ar}}{2} \left(z + \frac{2Q^2}{m_{ar}^2} \right) \right) , \quad (27)$$

$$p_c = \left(\frac{m_{ar}}{2} (1-z); -\sqrt{(1-z)Q^2 - \frac{Q^4}{m_{ar}^2}}, 0, \frac{m_{ar}}{2} \left(1-z - \frac{2Q^2}{m_{ar}^2} \right) \right) . \quad (28)$$

For simplicity we have here put the azimuthal angle $\varphi = 0$.

Note that

$$p_{\perp b,c}^2 = (1-z)Q^2 - \frac{Q^4}{m_{ar}^2} < (1-z)Q^2 = p_{\perp\text{evol}}^2 . \quad (29)$$

For small Q^2 values the two measures $p_{\perp b,c}^2$ and $p_{\perp\text{evol}}^2$ agree well, but with increasing Q^2 the $p_{\perp b,c}^2$ will eventually turn over and decrease again (for fixed z and m_{ar}). Simple inspection shows that the maximum $p_{\perp b,c}^2$ occurs for $p_{\parallel c} = 0$

and that the decreasing $p_{\perp b,c}^2$ corresponds to increasingly negative $p_{\parallel c}$. The drop of $p_{\perp b,c}^2$ thus is deceptive, and does not correspond to our intuitive picture of time ordering. Like for the FSR algorithm, $p_{\perp \text{evol}}^2$ therefore makes more sense than $p_{\perp b,c}^2$ as evolution variable, in spite of it not always having as simple a kinematics interpretation. One should note, however, that emissions with negative $p_{\parallel c}$ are more likely to come from radiation off the other incoming parton, where it is collinearly enhanced, so in practice the region of decreasing $p_{\perp b,c}^2$ is not so important.

The allowed range $z_{\min} < z < z_{\max}$ is from below constrained by $x_a = x_b/z < 1$, i.e. $z_{\min} = x_b$, and from above by $p_{\perp b,c}^2 > 0$, which gives

$$z_{\max} = 1 - \frac{p_{\perp \text{evol}}}{m_{br}} \left(\sqrt{1 + \frac{p_{\perp \text{evol}}^2}{4m_{br}^2} - \frac{p_{\perp \text{evol}}}{2m_{br}}} \right). \quad (30)$$

When the $a \rightarrow bc$ kinematics is constructed, the above equations for $p_{a,b,c,r}$ are not sufficient. One also needs to boost and rotate all the partons produced by the incoming b and r partons. The full procedure then reads

1. Go to the $b + r$ rest frame, with $p_{b,r} = (m_{br}/2) (1; 0, 0, \pm 1)$.
2. Rotate by a randomly selected azimuthal angle $-\varphi$.
3. Put the b off mass shell, $Q^2 = -m_b^2 = p_{\perp \text{evol}}^2/(1-z)$, while preserving the total $b + r$ four-momentum, i.e. $p_{b,r} = ((m_{br}^2 \mp Q^2)/2m_{br}; 0, 0, \pm(m_{br}^2 + Q^2)/2m_{br})$.
4. Construct the massless incoming p_a in this frame, and the outgoing c , from the requirements $(p_a + p_r)^2 = m_{br}^2/z$ and $p_c^2 = (p_a - p_b)^2 = m_c^2 (= 0)$, and with transverse momentum in the x direction [4].
5. Boost everything to the $a + r$ rest frame, and thereafter rotate in θ to have a moving along the $+z$ axis.
6. Finally rotate $+\varphi$ in azimuth, with the same φ is in point 2. This gives c a random φ distribution, while preserving the φ values of the $b + r$ daughters, up to recoil effects.

Apart from the change of evolution variable, the major difference relative to the old algorithm [4] is that kinematics is now constructed with the recoiler assumed massless, rather than only after it has been assigned a virtuality as well.

Currently a smooth merging with first-order matrix elements is only available for the production of $\gamma^*/Z^0/W^\pm$ [31] and $gg \rightarrow H^0$ (in the infinitely-heavy-top-mass limit). It turns out that the shower actually does a reasonable job of describing radiation also harder than the mass scale of the electroweak production process, i.e. the matrix-element reweighting factors are everywhere of the order of unity. Unless there are reasons to the contrary, for non-QCD processes it therefore makes sense to start the shower from a $p_{\perp \text{max}} = \sqrt{s}/2$. For a normal QCD process this would lead to doublecounting, since the shower

emissions could be harder than the original hard process, but this risk does not exist for particles like the Z^0 , which are not produced in the shower anyway.

Mass corrections

Quark mass effects are seldom crucial for ISR: nothing heavier than charm and bottom need be considered as beam constituents, unlike the multitude of new massive particles one could imagine for FSR. Here the mass effects are less trivial to handle, however, since we may get stuck in impossible corners of phase space.

To illustrate this, consider $g \rightarrow Q\bar{Q}$, where we let Q denote a generic heavy quark, charm or bottom. Then requiring the lightcone $p_{\perp\text{LC}}^2 = (1-z)Q^2 - zm_Q^2 > 0$, eq. (9) with $m_a = 0$, $m_c = m_Q$ and $Q^2 = -m_b^2$, implies $z < Q^2/(Q^2 + m_Q^2)$. Since $x_a = x_b/z < 1$ it follows that the Q parton density must vanish for $x > Q^2/(Q^2 + m_Q^2)$. Many parton density parameterizations assume vanishing Q density below $Q^2 = m_Q^2$ and massless evolution above it, and so do not obey the above constraint.

Actually, with our energy-sharing z definition, now slightly modified but still preserving $z = m_{br}^2/m_{ar}^2$, eq. (29) is generalized to

$$p_{\perp b,c}^2 = (1-z)Q^2 - \frac{Q^4}{m_{ar}^2} - m_Q^2 \left(z + \frac{Q^2}{m_{ar}^2} \right) = Q^2 - z \frac{(Q^2 + m_Q^2)(m_{br}^2 + Q^2)}{m_{br}^2}. \quad (31)$$

which implies the somewhat tighter constraint

$$x_b < z < \frac{Q^2}{Q^2 + m_Q^2} \frac{m_{br}^2}{m_{br}^2 + Q^2}. \quad (32)$$

For the backwards evolution of $g \rightarrow Q\bar{Q}$, the evolution variable is chosen to be

$$p_{\perp\text{evol}}^2 = (1-z)(Q^2 + m_Q^2) = m_Q^2 + p_{\perp\text{LC}}^2, \quad (33)$$

such that a threshold set at $p_{\perp\text{evol}}^2 = m_Q^2$ corresponds to $p_{\perp\text{LC}}^2 \rightarrow 0$. Thereby, the evolution scale $p_{\perp\text{evol}}^2$ may be used as argument for α_s and for parton densities, while the physical p_{\perp} will still populate the full phase space.

Writing the upper limit in terms of the evolution variable $p_{\perp\text{evol}}^2$ rather than Q^2 , one obtains the analogue of eq. (30),

$$z_{\text{max}} = 1 - \frac{p_{\perp\text{evol}}}{m_{br}} \frac{1}{1 - \frac{m_Q^2}{m_{br}^2} \left(1 + \frac{m_{br}^2}{p_{\perp\text{evol}}^2} \right)} \left(\sqrt{1 + \frac{p_{\perp\text{evol}}^2 - m_Q^2}{4m_{br}^2}} - \frac{p_{\perp\text{evol}}}{2m_{br}} \left(1 + \frac{m_Q^2}{p_{\perp\text{evol}}^2} \right) \right). \quad (34)$$

This expression would be rather cumbersome to deal with in practice, but is bounded from above,

$$z_{\max} < \frac{m_{br}(m_{br} - m_Q)}{m_{br}^2 + m_Q m_{br} - m_Q^2}, \quad (35)$$

which we make use of in the evolution.

Should a hard-scattering configuration be inconsistent with these constraints, it is rejected as unphysical. Should the shower end up in such a region during the backwards evolution, a new shower is generated. Even when no such disasters occur, the fact that the physically allowed z range is smaller than what has been assumed in standard parton density parameterizations implies that more heavy quarks can survive to the near-threshold region than ought to be the case. This could be amended by an ad hoc compensating weight factor in the splitting kernel, but currently we have not studied this further.

Another technical problem is that, when performing the backwards evolution, eq. (19), one needs to estimate from above the ratio of parton densities, in order for the veto algorithm to be applicable [6]. Normally, densities fall off with x (the exception being valence quarks, for which some extra consideration is required) and have a modest scale dependence, so that

$$\frac{x' f_a(x', p_{\perp \text{evol}}^2)}{x f_b(x, p_{\perp \text{evol}}^2)} < \frac{x f_a(x, p_{\perp \text{evol}}^2)}{x f_b(x, p_{\perp \text{evol}}^2)} \simeq \frac{x f_a(x, p_{\perp \text{max}}^2)}{x f_b(x, p_{\perp \text{max}}^2)}. \quad (36)$$

Now, however, the denominator $f_b = f_Q$ vanishes for $p_{\perp \text{evol}}^2 \rightarrow m_Q^2$, and so does not obey the above relation. Given that $f_Q(x, Q^2)$ increases roughly like $\ln(Q^2/m_Q^2)$, a reasonable alternative approximation is

$$\frac{x' f_g(x', p_{\perp \text{evol}}^2)}{x f_Q(x, p_{\perp \text{evol}}^2)} < \frac{x f_g(x, p_{\perp \text{evol}}^2)}{x f_Q(x, p_{\perp \text{evol}}^2)} \simeq \frac{\ln(p_{\perp \text{max}}^2/m_Q^2)}{\ln(p_{\perp \text{evol}}^2/m_Q^2)} \frac{x f_g(x, p_{\perp \text{max}}^2)}{x f_Q(x, p_{\perp \text{max}}^2)}. \quad (37)$$

The $1/\ln(p_{\perp \text{evol}}^2/m_Q^2)$ prefactor can be incorporated into the choice of the next trial emission, so that steps taken in $p_{\perp \text{evol}}^2$ get shorter and shorter as the threshold is approached, until a valid branching is found.

Finally, the $g \rightarrow Q\bar{Q}$ splitting function should be modified. The appropriate expressions may be identified by considering the collinear limit of relevant matrix elements. Neglecting overall factors, $g \rightarrow Q\bar{Q}$ is equivalent to $\gamma \rightarrow \mu^+\mu^-$ with massive muons. Considering the $t \rightarrow 0$ limit of processes such as $\gamma\nu_\mu \rightarrow \mu^-W^+$ and $\gamma\mu^- \rightarrow \mu^-H^0$, and letting $m_\mu^2/m_{W,H}^2 \rightarrow 0$, we thus obtain:

$$P_{g \rightarrow Q\bar{Q}}(z) = \frac{1}{2} \left(z^2 + (1-z)^2 + 2z(1-z) \frac{m_Q^2}{p_{\perp \text{evol}}^2} \right), \quad (38)$$

which approaches a flat 1/2 for $p_{\perp \text{evol}}^2 \rightarrow m_Q^2$.

Since $g \rightarrow Q\bar{Q}$ and $Q \rightarrow Qg$ compete in the backwards evolution of a heavy quark, the $p_{\perp\text{evol}}^2 = (1-z)(Q^2 + m_Q^2)$ of eq. (33) is used also here. The kinematics interpretation is now slightly different, however. The branching Q is forced to be massless, so the kinematics is in this case identical to that of a light-quark $q \rightarrow qg$ branching. However, since the massive $p_{\perp\text{evol}}^2$ is different from the massless one, the z limit expressed in terms of $p_{\perp\text{evol}}^2$ also becomes different from eq. (30):

$$z_{\text{max}} = 1 - \frac{p_{\perp\text{evol}}}{m_{br}} \frac{1}{1 - \frac{m_Q^2}{m_{br}^2}} \left(\sqrt{1 + \frac{p_{\perp\text{evol}}^2}{4m_{br}^2} \left(1 - \frac{m_Q^2}{p_{\perp\text{evol}}^2}\right)^2} - \frac{p_{\perp\text{evol}}}{2m_{br}} \left(1 + \frac{m_Q^2}{p_{\perp\text{evol}}^2}\right) \right). \quad (39)$$

As before, also the splitting kernel receives a mass correction. For $Q \rightarrow Qg$, this may be obtained by considering the equivalent processes $\mu^- \bar{\nu}_\mu \rightarrow \gamma W^-$ and $\mu^+ \mu^- \rightarrow \gamma H^0$ in the same limits as above, yielding:

$$P_{Q \rightarrow Qg}(z) = \frac{4}{3} \left(\frac{1+z^2}{1-z} - 2z(1-z) \frac{m_Q^2}{p_{\perp\text{evol}}^2} \right), \quad (40)$$

i.e. the mass correction here has the same form but the opposite sign as for $g \rightarrow Q\bar{Q}$.

Finally, in the branching $Q \rightarrow gQ$, a gluon is emitted by a heavy quark, which in its turn must come from a $g \rightarrow Q\bar{Q}$ branching. Thus both the Q and \bar{Q} must be put on the mass shell, which implies significant kinematical constraints. The process is rare, however, and currently we have not considered it further.

Algorithm tests

While a FSR algorithm can be tested in e^+e^- annihilation events, where only hadronization need be considered in addition, the busier environment in hadron colliders makes ISR algorithms more complicated to test. One of the few clean measurements is provided by the p_{\perp} spectrum of Z^0 bosons. This quantity has been studied for the new algorithm (without the inclusion of incoming heavy flavours) [32], with the conclusion that it there does at least as well as the old PYTHIA algorithm. This is not surprising since the two are not so very different, apart from the Q^2 vs. p_{\perp}^2 ordering issue.

Actually, below and around the $d\sigma/dp_{\perp Z}$ peak, at $p_{\perp Z} \approx 4$ GeV at the Tevatron, a difference would have been welcome, since the old algorithm requires an uncomfortably large primordial k_{\perp} of around 2 GeV to provide a decent fit. Unfortunately the new requires about the same. The number can

be reduced by using a larger Λ in the algorithms than that of the parton densities. Such a procedure can be motivated by noting that the actual evolution in a generator contains various kinematical and dynamical suppressions not found in the leading-log parton evolution equations [32]. A fit to the whole $d\sigma/dp_{\perp Z}$ spectrum in the peak region does not favour significant reductions of the primordial k_{\perp} , however. This might be viewed as indications for the need of physics beyond standard DGLAP [22].

2.6 Combining spacelike and timelike showers

The separation of ISR and FSR is not unambiguous: it is possible to shuffle contributions between the two, i.e. take fewer but longer steps in rapidity for the ISR and compensate that by more extensive FSR radiation off those ISR partons that are emitted [24]. In part, compensation mechanisms of this kind automatically occur: if ISR partons are more widely spaced then the colour dipoles spanned between them become larger and thereby the FSR is increased, at least to some extent.

We defer further studies of the optimal balance between the two, and for now pick a simple strategy:

- The initial-state shower is first handled in full. This provides a set of final-state partons, from the hard interactions and from the c partons of all $a \rightarrow bc$ branchings in the ISR chains.
- Each final-state parton is associated with a p_{\perp} scale at which it was formed, either the hard-scattering scale or the $p_{\perp\text{evol}}$ of the ISR evolution.
- Each coloured final-state parton is also connected to other final-state partons to form colour dipoles. Normally these dipole partners would also act as recoilers. Top decay has been mentioned as one example where this would not be the case, but such decays can be considered separately from the production processes studied here, and before the tops decay they can act both as radiators and recoilers. When a colour-singlet particle like the Z^0 is produced, there is a freedom to admit this as a recoiler, to the hardest parton emitted on either side of it, or to let those two partons act as each other's recoilers, just like they are colour-connected. For now we choose the latter strategy.
- The lowest- $p_{\perp\text{evol}}$ parton emitted on either side of the event is colour-connected to the beam remnant. A remnant does not radiate, but can act as recoiler; since the momentum transfer will predominantly be in the longitudinal direction, it will not give rise to any unphysical p_{\perp} kicks. The internal structure of the remnant then has to be resolved beforehand, since a small radiator+recoiler invariant mass implies a restricted phase space for emissions. Such a dependence of perturbative physics on nonperturbative assumptions may be a bit uncomfortable. As an option, we have studied a scenario without any emissions at all off this radiator+recoiler set. Since the

affected parton normally is a low- $p_{\perp\text{evol}}$ one, and the potential additional activity should occur below this already low scale, one would not expect large differences, and indeed this is confirmed by our studies.

- The issue of what to do with loose colour ends is more important if one intends to stop and restart the showers (both ISR and FSR) at large p_{\perp} scales, as in a (L)CKKW-style matching to higher-order matrix-element programs [19]. We therefore consider two alternatives for the FSR activity off the dipoles defined by the ISR branchings. In one, each parton of a dipole radiates with a maximum p_{\perp} scale set by its production p_{\perp} , phase space constraints permitting. In the other, the maximum radiation scale in a dipole is set by the smaller of the two endpoint parton production p_{\perp} values, i.e. a dipole does not radiate above the scale at which it is ‘formed’. Technically, the latter option offers the possibility to combine ISR and FSR emissions in one common sequence of decreasing p_{\perp} values, certainly a boon for matching procedures. The choice of maximum emission scale is not unique, since the shower language offers little guidance in the regions where several p_{\perp} values are of comparable magnitude. In this case, that would be the emission or not of a hard FSR parton off the harder of the ISR ones. Practical experience could tell which is preferable.
- For now, however, all ISR activity is finished before the system is evolved with the FSR algorithm, downwards in $p_{\perp\text{evol}}$. Initially only the hardest partons can therefore radiate, but as $p_{\perp\text{evol}}$ is reduced also more of the partons from the ISR cascades can radiate, below the respective scale at which they themselves or their dipole were produced, depending on the option used.

3 Interleaved Multiple Interactions

3.1 Multiple interactions

Our basic framework for multiple interactions is the one presented in ref. [1], which in turn builds on the work in ref. [2]. We refer the reader to these for details, and here only provide a very brief summary.

The basic formalism

The cross section for $2 \rightarrow 2$ QCD scatterings is dominated by t -channel gluon exchange and hence diverges roughly like dp_{\perp}^2/p_{\perp}^4 . Therefore the integrated interaction cross section above some $p_{\perp\text{min}}$ scale, $\sigma_{\text{int}}(p_{\perp\text{min}})$, exceeds the total inelastic nondiffractive cross section σ_{nd} when $p_{\perp\text{min}} \rightarrow 0$. The resolution of this apparently paradoxical situation probably comes in two steps.

Firstly, the interaction cross section is an inclusive number. Thus, if an event contains two interactions it counts twice in σ_{int} but only once in σ_{nd} , and so on for higher multiplicities. Thereby we may identify $\langle n \rangle(p_{\perp\text{min}}) =$

$\sigma_{\text{int}}(p_{\perp\text{min}})/\sigma_{\text{nd}}$ with the average number of interactions above $p_{\perp\text{min}}$ per inelastic nondiffractive event, and that number may well be above unity.

As a starting point we will assume that all hadronic collisions are equivalent, i.e. that there is no dependence on impact parameter, and that the different parton-parton interactions take place independently of each other, i.e. we disregard energy-momentum conservation effects. The number of interactions above $p_{\perp\text{min}}$ per event is then distributed according to a Poisson distribution with mean $\langle n \rangle$, $\mathcal{P}_n = \langle n \rangle^n \exp(-\langle n \rangle)/n!$.

Secondly, the incoming hadrons are colour singlet objects. Therefore, when the p_{\perp} of an exchanged gluon is made small and the transverse wavelength correspondingly large, the gluon can no longer resolve the individual colour charges, and the effective coupling is decreased. Note that perturbative QCD calculations are always performed assuming free incoming and outgoing quark and gluon states, rather than partons inside hadrons, and thus do not address this kind of nonperturbative screening effects.

The simplest solution to the second issue is to introduce a step function $\theta(p_{\perp} - p_{\perp\text{min}})$, such that the perturbative cross section is assumed to completely vanish below some $p_{\perp\text{min}}$ scale. Given the complexity of the nonperturbative physics involved, $p_{\perp\text{min}}$ cannot be calculated but has to be tuned to data. A more realistic alternative is to note that the jet cross section is divergent like $\alpha_s^2(p_{\perp}^2)/p_{\perp}^4$, and that therefore a factor

$$\frac{\alpha_s^2(p_{\perp 0}^2 + p_{\perp 1}^2)}{\alpha_s^2(p_{\perp}^2)} \frac{p_{\perp}^4}{(p_{\perp 0}^2 + p_{\perp}^2)^2} \quad (41)$$

would smoothly regularize the divergences, now with $p_{\perp 0}$ as the free parameter to be tuned to data. Later we will return to the issue of whether to do a similar replacement for the scale argument of parton densities.

In an event with several interactions, it is convenient to order them in p_{\perp} , as already discussed in the introduction. The generation of a sequence $\sqrt{s}/2 > p_{\perp 1} > p_{\perp 2} > \dots > p_{\perp n} > p_{\perp\text{min}}$ now becomes one of determining $p_{\perp} = p_{\perp i}$ from a known $p_{\perp i-1}$, according to the probability distribution

$$\frac{d\mathcal{P}}{dp_{\perp}} = \frac{1}{\sigma_{\text{nd}}} \frac{d\sigma}{dp_{\perp}} \exp \left[- \int_{p_{\perp}}^{p_{\perp i-1}} \frac{1}{\sigma_{\text{nd}}} \frac{d\sigma}{dp'_{\perp}} dp'_{\perp} \right]. \quad (42)$$

The exponential expression is the ‘form factor’ from the requirement that no interactions occur between $p_{\perp i-1}$ and $p_{\perp i}$, cf. the Sudakov form factor of parton showers.

More realistically, one should include the possibility that each collision also could be characterized by a varying impact parameter b . Within the classical framework we use here, b is to be thought of as a distance of closest approach, not as the Fourier transform of the momentum transfer. A small b value corresponds to a large overlap between the two colliding hadrons, and hence an

enhanced probability for multiple interactions. A large b , on the other hand, corresponds to a grazing collision, with a large probability that no parton-parton interactions at all take place.

Let $\mathcal{O}(b)$ denote the time-integrated matter overlap between the two incoming hadrons at impact parameter b . The combined selection of b and a set of scattering $p_{\perp i}$ values can be reduced to a combined choice of b and $p_{\perp 1}$, according to a generalization of eq. (42)

$$\frac{d\mathcal{P}}{dp_{\perp 1} d^2b} = \frac{\mathcal{O}(b)}{\langle \mathcal{O} \rangle} \frac{1}{\sigma_{\text{nd}}} \frac{d\sigma}{dp_{\perp}} \exp \left[-\frac{\mathcal{O}(b)}{\langle \mathcal{O} \rangle} \int_{p_{\perp}}^{\sqrt{s}/2} \frac{1}{\sigma_{\text{nd}}} \frac{d\sigma}{dp'_{\perp}} dp'_{\perp} \right]. \quad (43)$$

The subsequent interactions can be generated sequentially in falling p_{\perp} as before, with the only difference that $d\sigma/dp_{\perp}^2$ now is multiplied by $\mathcal{O}(b)/\langle \mathcal{O} \rangle$, where b is fixed at the value chosen above.

Correlated parton densities

Consider a hadron undergoing multiple interactions in a collision. Such an object should be described by multi-parton densities, giving the joint probability of simultaneously finding n partons with flavours f_1, \dots, f_n , carrying momentum fractions x_1, \dots, x_n inside the hadron, when probed by interactions at scales Q_1^2, \dots, Q_n^2 , in our case with the association $Q_i^2 = p_{\perp i}^2$. Having nowhere near sufficient experimental information to pin down such distributions, and wishing to make maximal use of the information that we *do* have, namely the standard one-parton-inclusive parton densities, we propose the following strategy.

The first and most trivial observation is that each interaction i removes a momentum fraction x_i from the hadron remnant. This momentum loss can be taken into account by assuming a simple scaling ansatz for the parton distributions, $f(x) \rightarrow f(x/X)/X$, where $X = 1 - \sum_{i=1}^n x_i$ is the momentum remaining in the beam hadron after the n first interactions. Effectively, the PDF's are simply 'squeezed' into the range $x \in [0, X]$.

Next, for a given hadron, the valence distribution of flavour f after n interactions, $q_{fvn}(x, Q^2)$, should integrate to the number N_{fvn} of valence quarks of flavour f remaining in the hadron remnant. This rule may be enforced by scaling the original distribution down, by the ratio of remaining to original valence quarks N_{fvn}/N_{fv0} , in addition to the x scaling mentioned above.

Also, when a sea quark is knocked out of a hadron, it must leave behind a corresponding antisea parton in the beam remnant. We call this a companion quark. In the perturbative approximation the sea quark q_s and its companion q_c come from a gluon branching $g \rightarrow q_s + q_c$ (it is implicit that if q_s is a quark, q_c is its antiquark). Starting from this perturbative ansatz, and neglecting other interactions and any subsequent perturbative evolution of the q_c , we obtain the q_c distribution from the probability that a sea quark q_s , carrying a momentum

fraction x_s , is produced by the branching of a gluon with momentum fraction y , so that the companion has a momentum fraction $x = y - x_s$,

$$q_c(x; x_s) \propto \int_0^1 g(y) P_{g \rightarrow q_s q_c}(z) \delta(x_s - zy) dz = \frac{g(x_s + x)}{x_s + x} P_{g \rightarrow q_s q_c} \left(\frac{x_s}{x_s + x} \right), \quad (44)$$

with $P_{g \rightarrow q_s q_c}$ the usual DGLAP gluon splitting kernel. A simple ansatz $g(x) \propto (1-x)^n/x$ is here used for the gluon. Normalizations are fixed so that a sea quark has exactly one companion.

Without any further change, the reduction of the valence distributions and the introduction of companion distributions, in the manner described above, would result in a violation of the total momentum sum rule, that the x -weighted parton densities should integrate to X : by removing a valence quark from the parton distributions we also remove a total amount of momentum corresponding to $\langle x_{fv} \rangle$, the average momentum fraction carried by a valence quark of flavour f , and by adding a companion distribution we add an analogously defined momentum fraction. To ensure that the momentum sum rule is still respected, we assume that the sea and gluon normalizations fluctuate up when a valence distribution is reduced and down when a companion distribution is added, by a multiplicative factor. The requirement of a physical x range is of course still maintained by ‘squeezing’ all distributions into the interval $x \in [0, X]$.

After the perturbative interactions have taken each their fraction of longitudinal momentum, the remaining momentum is to be shared between the beam remnant partons. Here, valence quarks receive an x picked at random according to a small- Q^2 valence-like parton density, while sea quarks must be companions of one of the initiator quarks, and hence should have an x picked according to the $q_c(x; x_s)$ distribution introduced above. In the rare case that no valence quarks remain and no sea quarks need be added for flavour conservation, the beam remnant is represented by a gluon, carrying all of the beam remnant longitudinal momentum.

Further aspects of the model include the possible formation of composite objects in the beam remnants (e.g. diquarks) and the addition of non-zero primordial k_\perp values to the parton shower initiators. Especially the latter introduces some complications, to obtain consistent kinematics. More complete descriptions may be found in [1, 20].

Colour correlations

The initial state of a baryon may be represented by three valence quarks, connected antisymmetrically in colour via a central junction, which acts as a switchyard for the colour flow and carries the net baryon number.

The colour-space evolution of this state into the initiator and remnant partons actually found in a given event is not predicted by perturbation the-

ory, but is crucial in determining how the system hadronizes; in the Lund string model [8], two colour-connected final-state partons together define a string piece, which hadronizes by successive non-perturbative breakups along the string. Thus, the colour flow of an event determines the topology of the hadronizing strings, and consequently where and how many hadrons will be produced. The question can essentially be reduced to one of choosing a fictitious sequence of gluon emissions off the initial valence topology, since sea quarks together with their companion partners are associated with parent gluons, by construction.

The simplest solution is to assume that gluons are attached to the initial quark lines in a random order. If so, the junction of an incoming baryon would rarely be colour-connected directly to two valence quarks in the beam remnant, and the initial-state baryon number would be able to migrate to large p_{\perp} and small x_F values. While such a mechanism should be present, there are reasons to believe that a purely random attachment exaggerates the migration effects. Hence a free parameter is introduced to suppress gluon attachments onto colour lines that lie entirely within the remnant.

This still does not determine the order in which gluons are attached to the colour line between a valence quark and the junction. We consider a few different possibilities: 1) random, 2) gluons are ordered according to the rapidity of the hard scattering subsystem they are associated with, and 3) gluons are ordered so as to give rise to the smallest possible total string lengths in the final state. The two latter possibilities correspond to a tendency of nature to minimize the total potential energy of the system, i.e. the string length. Empirically such a tendency among the strings formed by multiple interactions is supported e.g. by the observed rapid increase of $\langle p_{\perp} \rangle$ with n_{charged} [33]

It appears, however, that a string minimization in the initial state is not enough, and that also the colours inside the initial-state cascades and hard interactions may be nontrivially correlated. Currently this is handled by a reassignment among a fraction of the colours in the final state, chosen so as to reduce the total string length.

3.2 Multiple interactions and initial-state radiation

Each multiple interaction is associated with its set of initial- and final-state radiation. We have already argued that, to a good approximation, the addition of FSR can be deferred until after ISR and MI have been considered in full. Specifically, FSR does not modify the total amount of energy carried by perturbatively defined partons, it only redistributes that energy among more partons. By contrast, both the addition of a further ISR branching and the addition of a further interaction implies more perturbative energy, taken from the limited beam-remnants reservoir. These two mechanisms therefore are in direct competition with each other.

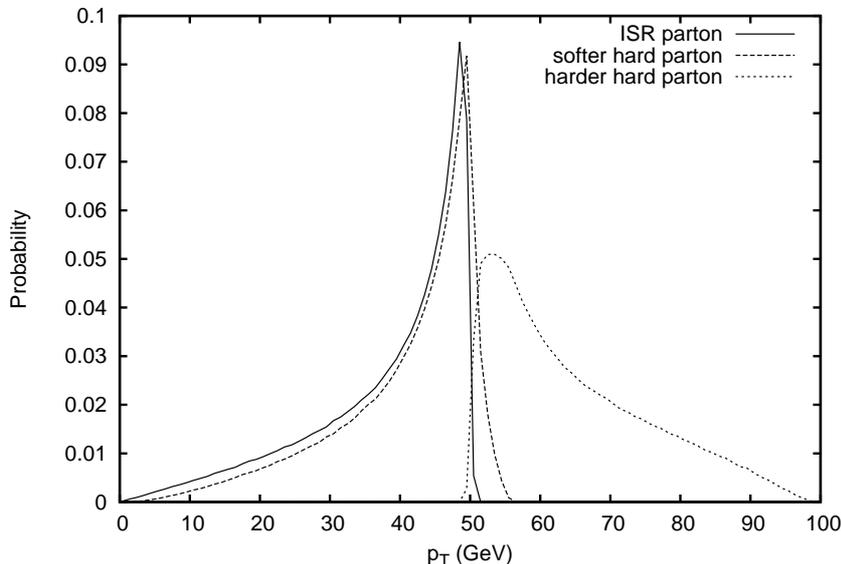


Figure 4: Parton p_{\perp} spectra when 2-parton events of a fixed $p_{\perp} = 50$ GeV, for an 1800 GeV $p\bar{p}$ collider, are modified by a single ISR branching with $p_{\perp\text{evol}} = 50$ GeV, using CTEQ5L parton distributions and the standard DGLAP splitting kernels. Owing to $p_{\perp\text{evol}} \neq p_{\perp}$, the parton emitted at the ISR branching has a tail to p_{\perp} values well below 50 GeV. However, this spectrum is comparable with the lower- p_{\perp} of the two hard-scattering partons, after the recoil from the ISR has been taken into account, so there is a certain symmetry if it all is viewed as a $2 \rightarrow 3$ process.

We have advocated for p_{\perp} as a convenient ordering variable, with smaller p_{\perp} values corresponding to ‘later times’. The p_{\perp} measure used for MI fills a similar function as the $p_{\perp\text{evol}}$ variable used for ISR, such that the two can be viewed as measuring the same kind of ‘time ordering’. To wit, kinematically $p_{\perp\text{evol}}$ agrees well with the standard p_{\perp} , except in the corner of high virtualities, where there is little multiple activity anyway. An example of this mapping is shown in Fig. 4. Further, the generation of a new interaction, eq. (42) (or its extension to varying impact parameters), can be viewed as an evolution downwards in a $p_{\perp\text{evol}} = p_{\perp}$, in a similar form-factor formalism as for the backwards evolution of ISR.

Starting from a hard interaction, a common sequence of subsequent evolution steps — interactions and branchings mixed — can therefore be found. Assuming that the latest step occurred at some $p_{\perp i-1}$ scale, this sets the maximum $p_{\perp\text{max}} = p_{\perp i-1}$ for the continued evolution. What can happen next is

then either a new interaction or a new ISR branching on one of the two incoming sides in one of the existing interactions. The probability distribution for $p_{\perp} = p_{\perp i}$ is given by

$$\frac{d\mathcal{P}}{dp_{\perp}} = \left(\frac{d\mathcal{P}_{\text{MI}}}{dp_{\perp}} + \sum \frac{d\mathcal{P}_{\text{ISR}}}{dp_{\perp}} \right) \exp \left(- \int_{p_{\perp}}^{p_{\perp i-1}} \left(\frac{d\mathcal{P}_{\text{MI}}}{dp'_{\perp}} + \sum \frac{d\mathcal{P}_{\text{ISR}}}{dp'_{\perp}} \right) dp'_{\perp} \right) \quad (45)$$

in simplified notation. Technically, the $p_{\perp i}$ can be found by selecting a new trial interaction according to $d\mathcal{P}_{\text{MI}} \exp(-\int d\mathcal{P}_{\text{MI}})$, and a trial ISR branching in each of the possible places according to $d\mathcal{P}_{\text{ISR}} \exp(-\int d\mathcal{P}_{\text{ISR}})$. The one of all of these possibilities that occurs at the largest p_{\perp} preempts the others, and is allowed to be realized. The whole process is iterated, until a lower cutoff is reached, below which no further interactions or branchings are allowed.

If there were no momentum constraints linking the different subsystems, it is easy to see that such an interleaved evolution actually is equivalent to considering the ISR of each interaction in full before moving on to the next interaction. Competition is introduced via the correlated parton densities already discussed. Thus distributions are squeezed to be nonvanishing in a range $x \in [0, X]$, where $X < 1$ represents the fraction of the original beam remnant momentum still available for an interaction or branching. When a trial n 'th interaction is considered, $X = 1 - \sum_{i=1}^{n-1} x_i$, where the sum runs over all the already existing interactions. The x_i are the respective momentum fractions of the ISR shower initiators at the current resolution scale, i.e., an x_i is increased each time an ISR branching is backwards-constructed on an incoming parton leg. Similarly, the flavour content is modified to take into account the partons already extracted by the $n - 1$ previous interactions, including the effects of ISR branchings. When instead a trial shower branching is considered, the X sum excludes the interaction under consideration, since this energy *is* at the disposal of the interaction, and similarly for the flavour content.

We have already discussed the choice of $p_{\perp \text{max}}$ scale for ISR showers, and that now generalizes. Thus, for minimum-bias QCD events the full phase space is allowed, while the p_{\perp} scale of a QCD hard process sets the maximum for the continued evolution, in order not to doublecount. When the hard process represents a possibility not present in the MI/ISR machinery — production of Z^0 , top, or supersymmetry, say — there is no risk of doublecounting, and again the full (remaining) phase space is available.

There is also the matter of a lower $p_{\perp \text{min}}$ scale. Customarily such scales are chosen separately for ISR and MI, and typically lower for the former than the latter. Both cutoffs are related to the resolution of the incoming hadronic wave function, however, and in the current formalism ISR and MI are interleaved, so it makes sense to use the same regularization procedure. Therefore also the branching probability is smoothly turned off at a $p_{\perp 0}$ scale, like for MI, by a factor the square root of eq. (41), since only one vertex is involved in a shower branching relative to the two of a hard process. Thus the

$\alpha_s(p_{\perp\text{evol}}^2) dp_{\perp\text{evol}}^2/p_{\perp\text{evol}}^2$ divergence is tamed to $\alpha_s(p_{\perp 0}^2 + p_{\perp\text{evol}}^2) dp_{\perp\text{evol}}^2/(p_{\perp 0}^2 + p_{\perp\text{evol}}^2)$. The scale of parton densities in ISR and MI alike is maintained at $p_{\perp\text{evol}}^2$, however, the argument being that the actual evolution of the partonic content is given by standard DGLAP evolution, and that it is only when this content is to be resolved that a dampening is to be imposed. This also has the boon that flavour thresholds appear where they are expected.

The cutoff for FSR still kept separate and lower, since that scale deals with the matching between perturbative physics and the nonperturbative hadronization at long time scales, and so has a somewhat different function.

4 Some First Results

4.1 Simple tunes

In this section, some first tests of the new framework are presented. We compare Tune A [34] of the old multiple interactions scenario [2] and the ‘‘Rap’’ tune of [1] with three rough ‘tunes’ of the new framework. These preliminary new tunes all take the parameters of the ‘‘Rap’’ model as a starting point:

- A matter overlap profile proportional to $\exp(-b^{1.8})$, where b is the impact parameter.
- Rapidity-ordered initial-state colour connections.
- Shower initiator attachments between two partons both in the beam remnant are suppressed by a factor 0.01 relative to others.
- Only valence quarks are allowed to participate in the formation of diquarks in the beam remnants, and these diquarks are then assumed to acquire total x values twice as large as the naive sum of x values of their constituents.
- As for Tune A, the regularization scale $p_{\perp 0}$ is given at a reference cm energy of 1800 GeV, with an energy rescaling proportional to $E_{\text{cm}}^{1/4}$.

These choices have been made for convenience, to keep down the number of free parameters to be tuned. Very likely, an improved agreement with data can be obtained by relaxing this, e.g. by varying the matter overlap profile. We also have indications that the energy dependence of $p_{\perp 0}$ may be smaller than in Tune A but, since we only show comparisons at 1.8–1.96 TeV, this will be of no importance here.

In addition, the three new tunes differ in the parameters listed in Table 1, which also show the resulting average numbers of interactions, and ISR and FSR branchings for each model in a ‘minimum-bias’ sample of inelastic non-diffractive events. One may view ‘‘High FSR’’ as our preferred new scenario, with ‘‘Sharp ISR’’ and ‘‘Low FSR’’ representing two variations, as a check of the sensitivity to some key assumptions.

The ‘‘Sharp ISR’’ model uses a threshold regularization of the ISR evolution (at $p_{\perp\text{evol}} = 1$ GeV), similarly to the old models, rather than the smooth

Kind	Model name	ISR cutoff	FSR scale	$p_{\perp 0}$ [GeV]	F' (F)	$\langle n_{\text{INT}} \rangle$	$\langle n_{\text{ISR}} \rangle$	$\langle n_{\text{FSR}} \rangle$
old	Tune A	sharp	–	2.00	–	5.8	2.0	3.6
”	Rap	sharp	–	2.40	(0.55)	3.6	4.4	5.5
new	Sharp ISR	sharp	radiator	2.70	1.9	1.8	3.9	15.9
”	Low FSR	smooth	lowest	2.30	0.8	2.9	2.2	9.2
”	High FSR	smooth	radiator	2.50	1.3	2.4	1.7	14.0

Table 1: The parameters distinguishing the new tunes, compared to Tune A and the “Rap” model where meaningful. Also shown are the mean numbers of interactions, $\langle n_{\text{INT}} \rangle$ (including the hardest), ISR branchings, $\langle n_{\text{ISR}} \rangle$, and FSR branchings, $\langle n_{\text{FSR}} \rangle$, for each model.

dampening, (the square root of) eq. (41), used in the other new tunes. (The multiple interactions cross sections are regularized by eq. (41) in all cases.) In addition, both the “High FSR” and the “Sharp ISR” tunes let the maximum scale for final-state emissions off a given parton be determined by the $p_{\perp \text{evol}}$ of that parton, while for the “Low FSR” tune the scale is given by whichever has the lowest $p_{\perp \text{evol}}$ of the two partons spanning the radiating dipole.

The parameter F' controls the strength of colour reconnections in the final state. Essentially, this is a fudge parameter, required in the new framework in order to approximately reproduce the effect of the rather extreme parameter settings controlling the final-state colour correlations between different scatterings in Tune A. We still have not penetrated to the details of the underlying mechanism here, i.e. why data seem to prefer such an extreme behaviour, hence the appearance of effective parameters controlling these correlations in both types of models. F' has a slightly different meaning than F of the “Rap” model, as follows. In [1], the colour reconnections were performed after all the perturbative activity had been generated, including final-state radiation. In the new framework, the colour reconnections are performed *before* the final-state showers, since *a priori* we believe it is mostly a lack of correlation in the initial-state colour flows that we are trying to make up for by this procedure.

The tunes have been produced by adjusting $p_{\perp 0}$ and F' so as to simultaneously describe the Tune A charged multiplicity and $\langle p_{\perp} \rangle (n_{\text{ch}})$ distributions as well as possible, since these in turn give good fits to Tevatron data. Results are shown in Figs. 5 and 6.

While the multiplicity distributions have been brought into fair agreement with each other, the Tune A $\langle p_{\perp} \rangle (n_{\text{ch}})$ is very difficult to duplicate in the new framework. This problem was also present for the models presented in [1]. Our interpretation is that this particular distribution is highly sensitive to the colour correlations, and we have so far been unsuccessful in identifying a physics mechanism that could explain the rather extreme correlations that are present

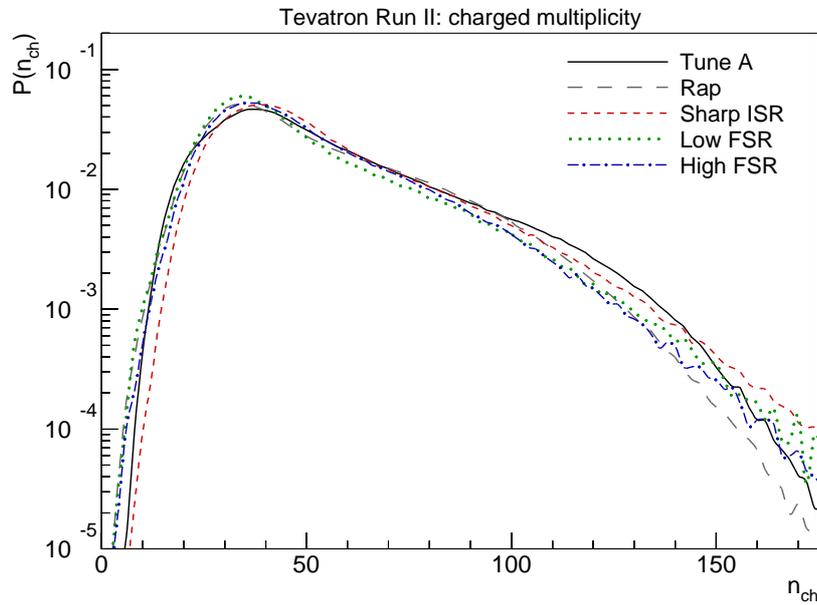


Figure 5: Charged multiplicity distributions, for 1.96 TeV $p\bar{p}$ minimum-bias events.

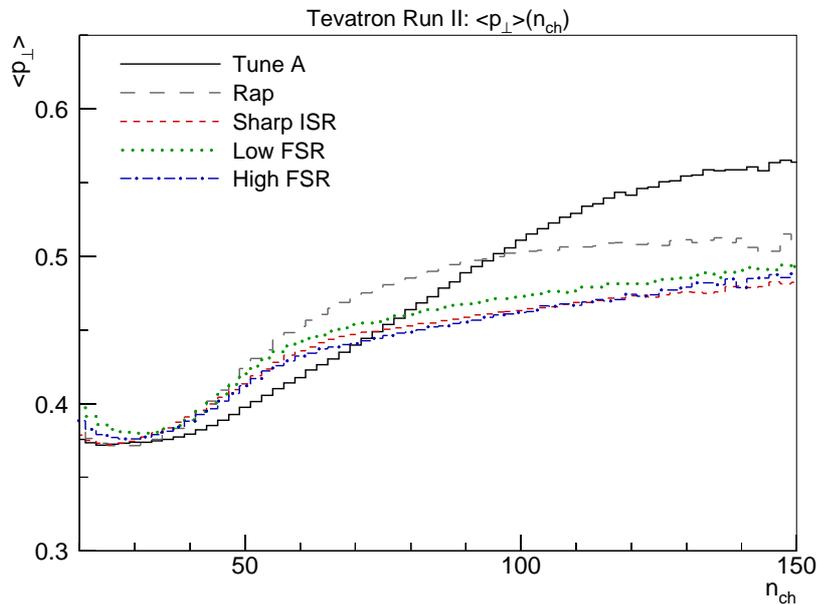


Figure 6: Average p_{\perp} as a function of charged multiplicity, $\langle p_{\perp} \rangle(n_{ch})$, for 1.96 TeV $p\bar{p}$ minimum-bias events. Note that the origo of the plot is *not* at (0,0).

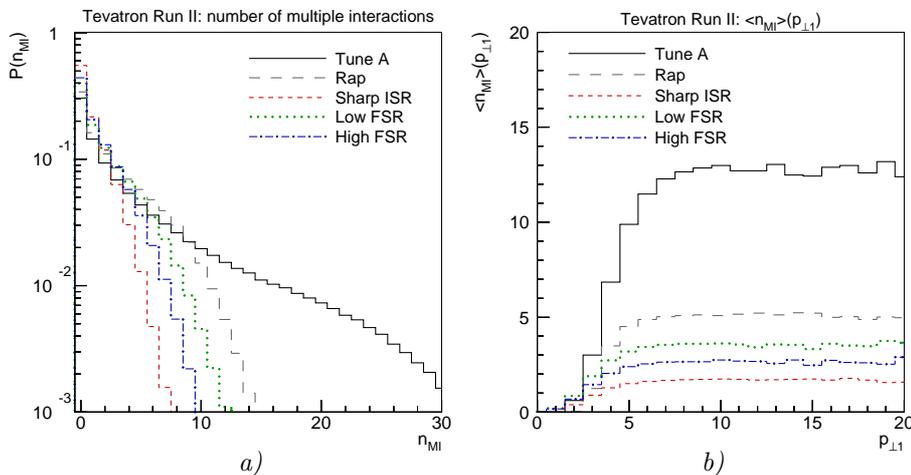


Figure 7: *a)* Number of multiple interactions (in addition to the hardest one) and *b)* the average number of additional interactions as a function of the p_{\perp} of the hardest interaction, both for 1.96 TeV $p\bar{p}$ minimum-bias events.

in Tune A. Since data seems to be in fair agreement with Tune A here, the bottom line is that *some* kind of more or less soft colour correlations working *between* the scattering chains is likely to be present, beyond what our primitive fudge parameters F and F' are capable of describing at this point.

4.2 Event activity

We now take a closer look at the relative proportions of the MI, ISR, and FSR make-up of minimum-bias events, for the models in Table 1. Firstly, the number of multiple interactions (excluding the hardest) is shown in Fig. 7a, and the dependence of the average number of extra interactions on the p_{\perp} of the hardest interaction in Fig. 7b. The relatively low $p_{\perp 0}$ and slightly more peaked matter distribution of Tune A gives a tail towards very large multiplicities which is substantially reduced both in the new models and in the Rap tune. Surprisingly, the Low FSR scenario lies somewhat below the Rap model, even though the latter has a higher $p_{\perp 0}$ scale. A sanity check is to switch off ISR and then compare the four models with the same matter overlap. Without the ISR evolution competing for phase space, the n_{MI} distribution then looks as would be expected, with the lower $p_{\perp 0}$ scenario exhibiting the broadest distribution. Thus, the ISR branchings ‘eat up’ phase space more quickly in the new framework than before, leaving less room for multiple interactions. This conclusion is verified in Fig. 8, which compares the distribution of z values for the first, i.e. hardest, ISR branching in an event. The soft-gluon enhancement

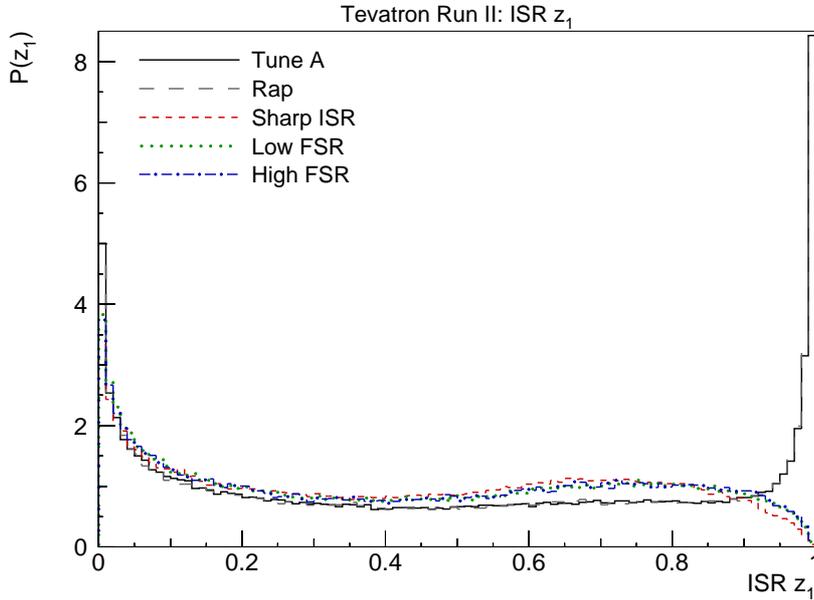


Figure 8: z distributions for the first ISR branching, z_1 , in 1.96 TeV $p\bar{p}$ minimum-bias events.

of ISR near $z = 1$ in the old models is absent in the new ones! This comes from the use of an evolution variable $p_{\perp\text{evol}}^2 = (1 - z)Q^2$ in the latter ones, which favours larger $1 - z$ in a branching than an evolution in Q^2 , cf. the z_{max} expression in eq. (30).

In analogy with Fig. 7, the multiplicities of ISR and FSR branchings are depicted in Figs. 9 and 10, respectively. For ISR as well as for FSR, Tune A has by far the narrowest distributions, since only the hardest interactions are associated with parton showers. Concentrating on the ISR distribution, Fig. 9, again the Rap model exhibits a very broad distribution, together with the Sharp ISR model. This behaviour is characteristic of the threshold regularization of the ISR cascade employed in these models, which gives a larger number of fairly soft emissions than the smoothly regularized models, Low and High FSR. Also note that the smaller number of branchings in these models partly is compensated by the larger $\langle 1 - z \rangle$ for the branchings that do occur.

The large number of FSR branchings, Fig. 10, is related to the use of a very small cutoff here, of the order of $p_{\perp\text{min}} = 0.5$ GeV, and so it cannot be compared directly with the MI and ISR multiplicities. The new models clearly have much broader FSR distributions than both Tune A and the Rap model. As one would expect from the choice of maximum scale of emission, the Low

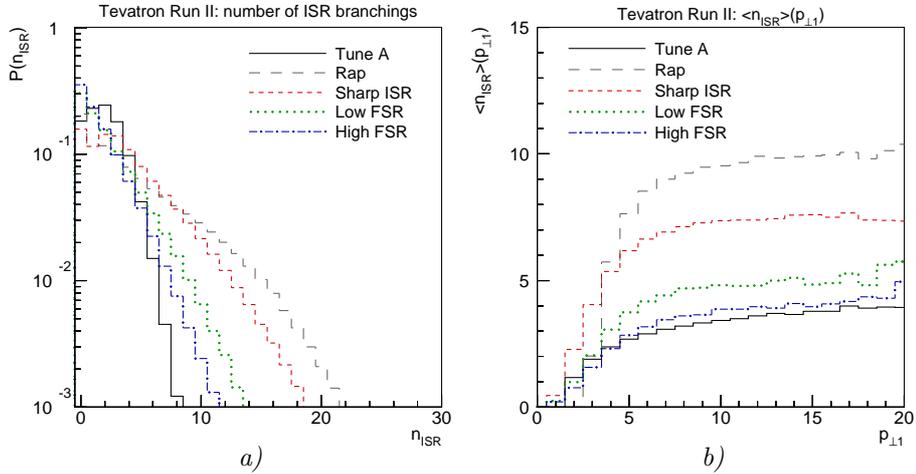


Figure 9: *a)* Number of ISR branchings and *b)* the average number of ISR branchings as a function of the p_{\perp} of the hardest interaction, both for 1.96 TeV $p\bar{p}$ minimum-bias events.

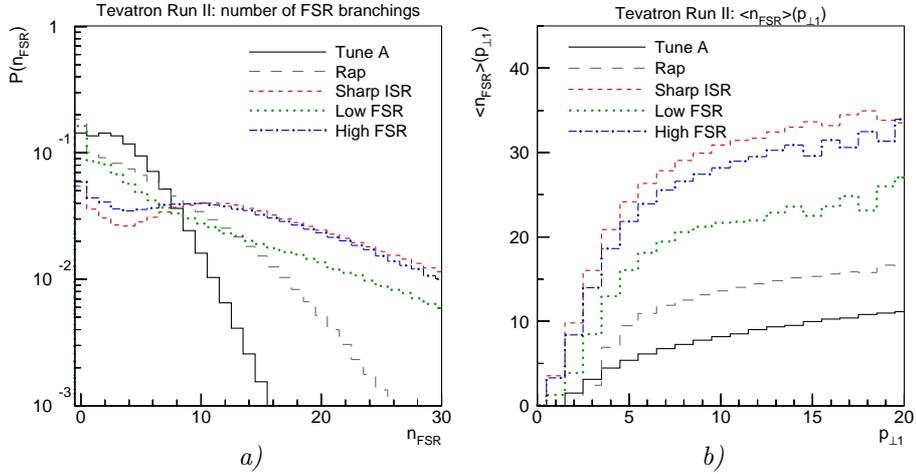


Figure 10: *a)* Number of FSR branchings and *b)* the average number of FSR branchings as a function of the p_{\perp} of the hardest interaction, both for 1.96 TeV $p\bar{p}$ minimum-bias events.

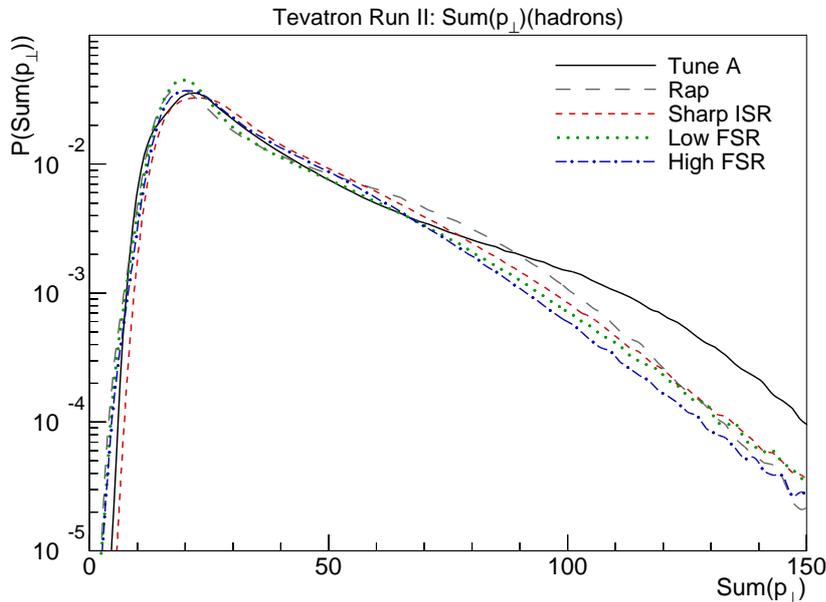


Figure 11: Total p_{\perp} sum for hadrons in 1.96 TeV $p\bar{p}$ minimum-bias events.

FSR model is the narrowest of the new models. We also recall that there is a built-in compensation mechanism: if the number of ISR branchings is reduced then, other things being the same, this results in fewer but larger dipoles that therefore can radiate more. Although the old and new shower algorithms do not allow a straightforward comparison, the difference between Rap and the new models is at least consistent with such a partial compensation.

Returning now to observable distributions, the fact that less p_{\perp} is kicked into events with large multiplicities in the new frameworks, cf. Fig. 6, while the multiplicity distributions are similar, also implies that there should be fewer events with large total E_{\perp} than in Tune A. This is corroborated by Fig. 11, which shows the scalar sum of hadron p_{\perp} values in 1.9 TeV $p\bar{p}$ minimum-bias events. Both the new models and the Rap model have noticeably fewer events in the region above ~ 100 GeV than does Tune A.

In addition, Fig. 12 shows that the pseudorapidity distribution has become narrower, i.e. the particle production has become more central. The normalization differences are in this context not very interesting, arising from small differences in the average charged multiplicity of the tunes.

However, these difference do not have a large impact on most other observables. Thus e.g. the minijet rates and charged hadron p_{\perp} distributions in Figs. 13 and 14 are hardly distinguishable between Tune A and the new

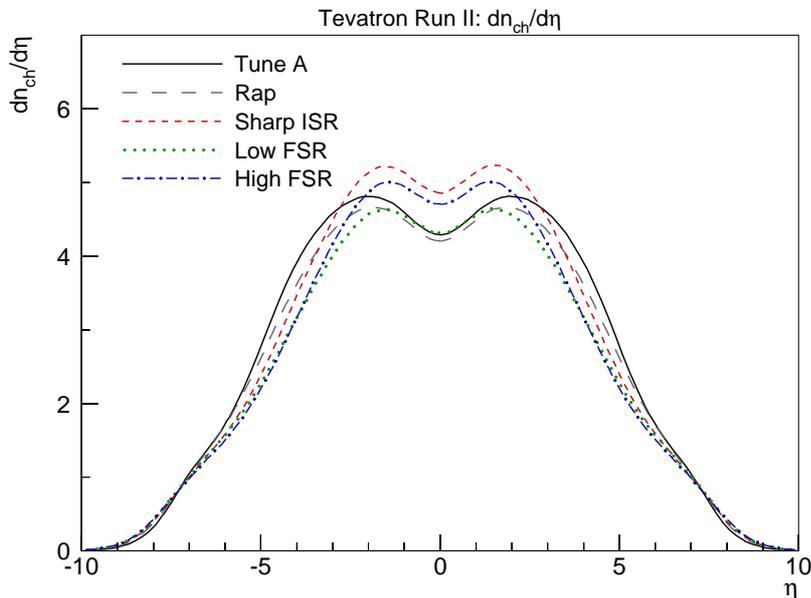


Figure 12: Charged multiplicity as a function of pseudorapidity η in 1.96 TeV $p\bar{p}$ minimum-bias events.

models. The minijet E_{\perp} spectrum, defined by a simple cone algorithm with a cone radius of $R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} = 0.7$ and an $E_{\perp\text{min}} = 5$ GeV, which was slightly softer in the Rap model than in Tune A, has become slightly harder. On the other hand, the charged hadron p_{\perp} spectrum, which was slightly harder in the Rap model than in Tune A, has dropped back down fairly close to the Tune A level.

4.3 Jet events and profiles

Complementary to the above are studies of events with hard jets and their properties. As an example of this, we have considered 1.96 TeV $p\bar{p}$ events where the hardest $2 \rightarrow 2$ interaction has a $p_{\perp\text{hard}} > 100$ GeV, without any further requirements. The charged multiplicity distribution of such events is shown in Fig. 15, and their pseudorapidity distribution in Fig. 16. Given that the models have been tuned to each other exclusively for a minimum-bias event sample, the differences are less than could have been expected. We note a clear difference at mid-rapidities, however, where Tune A shows more activity than any of the newer scenarios, cf. also Fig. 12. This is likely to be related to the way strings are connected from the central interactions to the beam remnants.

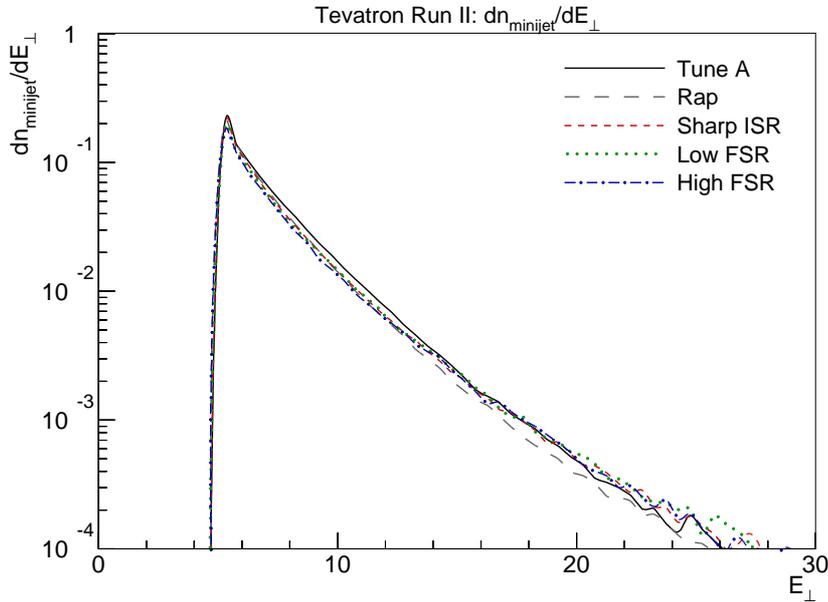


Figure 13: Minijet rates, for 1.96 TeV $p\bar{p}$ minimum-bias events.

The jet multiplicity in these events, obtained by a combination of MI, ISR and FSR activity, is shown in Fig. 17. The Low FSR scenario stands out by having significantly less jet activity than any of the other ones, clearly indicating the impact of the reduced FSR in these events. The other rates come surprisingly close, given that both the ISR and the FSR algorithms are quite different between the old and the new scenarios. At high jet multiplicities the new ones are somewhat above the older ones.

Next we study the properties of the jets produced. Since the two hardest jets both arise already as a consequence of the hard interaction, they have similar properties, while further jets are related to the additional activity and thus internally similar. Therefore only results for the hardest and (when present) third hardest jet are shown here. The respective jet E_{\perp} spectra are shown in Fig. 18. The hardest jet is harder in all the three new scenarios than in the two old ones, while the third and subsequent ones are more similar. Again, given the changed ISR and FSR algorithms, the similarities for the third jet are more surprising than the differences for the first. Notably, the lower jet activity in the Low FSR scenario is not reflected in a reduced tail out to high- E_{\perp} third jets.

The energy flow inside a jet can be plotted as a function of the distance r away from the center of the jet, or better as a function of r/R . Such profiles

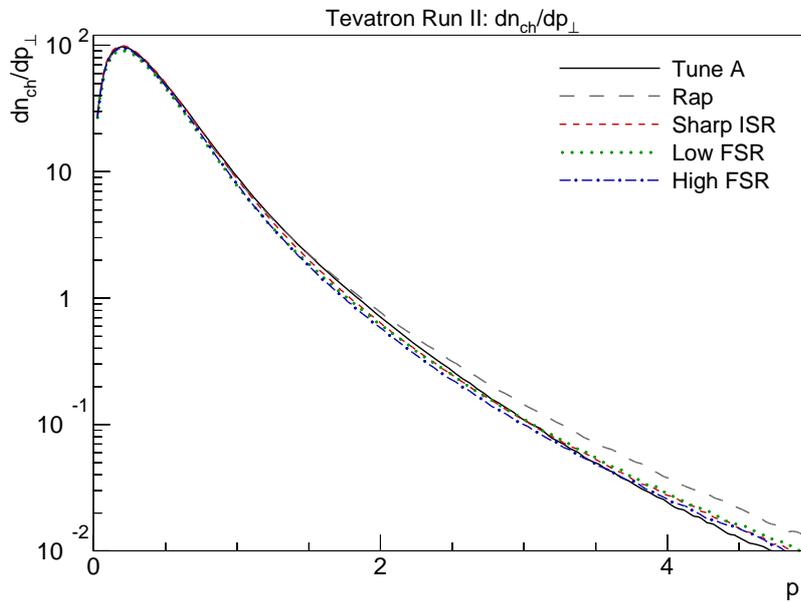


Figure 14: Charged hadron p_{\perp} spectra in 1.96 TeV $p\bar{p}$ minimum-bias events. The normalization corresponds the total average charged multiplicity.

are shown for the hardest and third jet in Fig. 19. For the hardest jet, again Low FSR stand out by producing narrower jets, while for the third Rap is even more narrow. Generally, the differences are small, however.

Turning to charged multiplicity distributions inside jets, the Rap scenario tends to have the least, and the High FSR and Sharp ISR the most. This is illustrated in Fig. 20a for the hardest jet, but the same pattern repeats also for the softer one. Comparing with the total charged multiplicity of these events, Fig. 15 above, which does not show the same pattern, we conclude that the balance between activity inside and outside the identified jets differs, possibly reflecting the amount of softer jet activity.

By contrast, the charged particle number jet profile follows the same pattern as observed above for the E_{\perp} profile. That is, Low FSR gives the most narrow hardest jet, Fig. 20b, while Rap gives the most narrow third jet, not shown.

In summary, differences are smaller than might have been guessed, considering the changes especially in the ISR and FSR algorithms. Specifically, with the new algorithms the upper scale $p_{\perp\max}$ for ISR and FSR evolution is unambiguously set by the $p_{\perp\text{hard}}$ of the hard interaction, while the older ones did involve an ambiguous choice of a $Q_{\max}^2 = 4p_{\perp\text{hard}}^2$, intended roughly to give p_{\perp} ordering, but not in the guaranteed sense of the new algorithms.

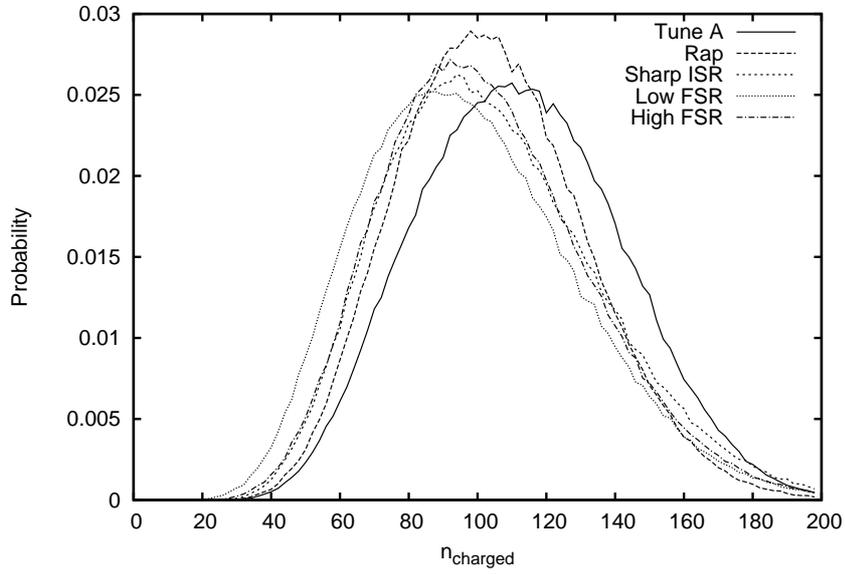


Figure 15: Charged multiplicity distribution for 1.96 TeV $p\bar{p}$ events with $p_{\perp\text{hard}} > 100$ GeV.

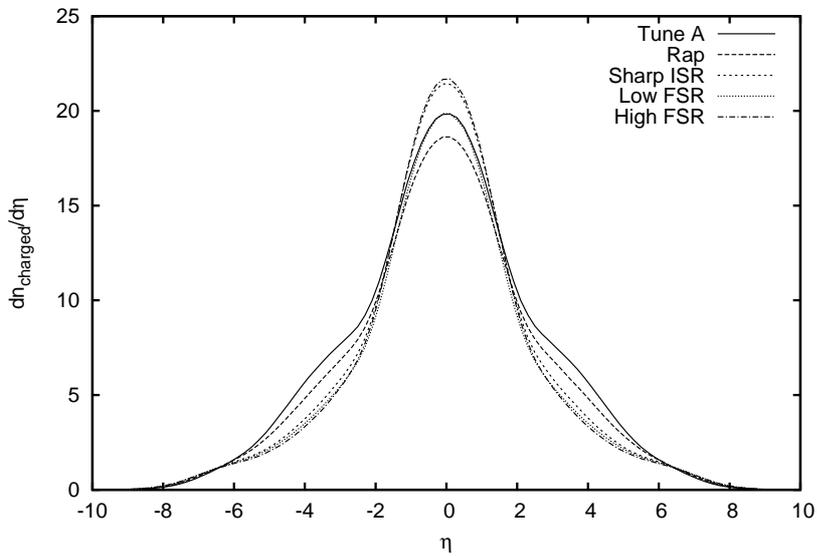


Figure 16: Pseudorapidity distribution for 1.96 TeV $p\bar{p}$ events with $p_{\perp\text{hard}} > 100$ GeV.

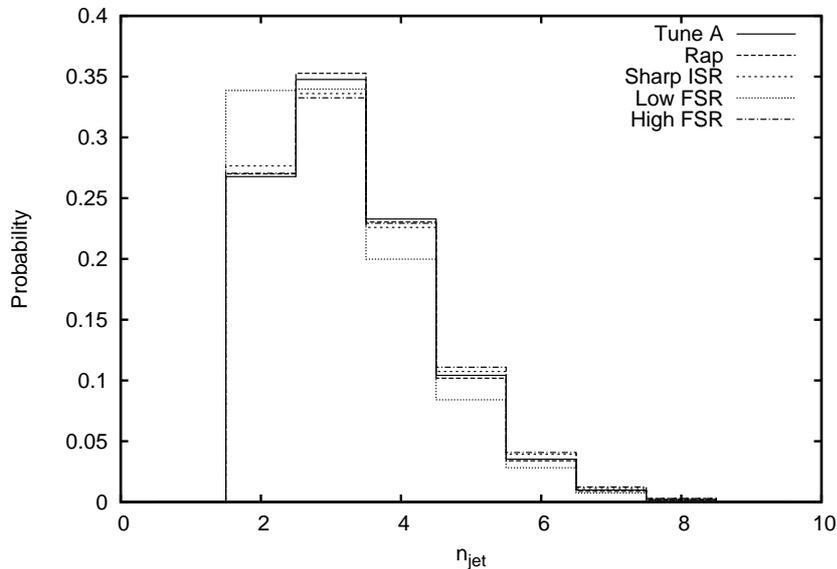


Figure 17: Jet multiplicity distribution for 1.96 TeV $p\bar{p}$ events with $p_{\perp\text{hard}} > 100$ GeV, using a cone clustering algorithm with $R = 0.7$ and $E_{\perp\text{min}} = 10$ GeV.

4.4 Z^0 production

A slightly different test is to study the p_{\perp} spectrum of high-mass dileptons coming from the decay of a γ^*/Z^0 . We can here compare with the CDF $p_{\perp Z}$ spectrum at 1.8 TeV [35], normalizing the curves to the experimental integrated cross section, Fig. 21.

Since the high- $p_{\perp Z}$ behaviour is constrained by our use of first-order matrix-element corrections [31], it is not surprising that differences here are small. That the three new scenarios are above the two older ones presumably is a consequence of the different treatment of FSR, which does not at all influence $p_{\perp Z}$ in the new models, while the p_{\perp} of an ISR branching is reduced by FSR in the older ones. This is a degree of freedom that could be studied further when FSR is interleaved with MI and ISR.

More interesting is the improvement in the low- p_{\perp} region, similarly to what has been found earlier [32], in a study of the new ISR algorithm without any MI. However, note that we in all cases make use of a Gaussian primordial k_{\perp} with a 2 GeV width (thus deviating from the pure Tune A, where it is kept at the PYTHIA default of 1 GeV). The implementation of this k_{\perp} is more complex with the new beam-remnant implementation of [1], and e.g. could depend on the number of multiple interactions, but actually the distributions turn out

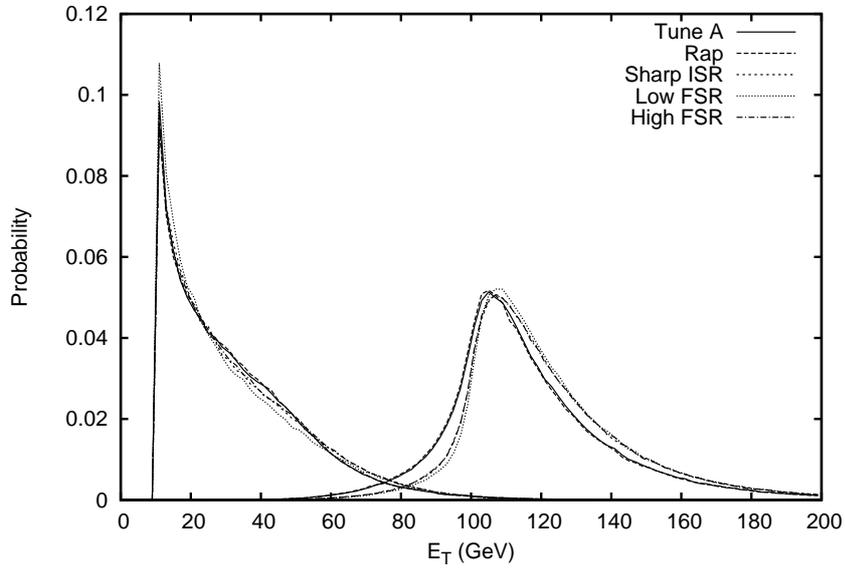


Figure 18: E_{\perp} spectra for the hardest and third hardest jet in 1.96 TeV $p\bar{p}$ events with $p_{\perp\text{hard}} > 100$ GeV.

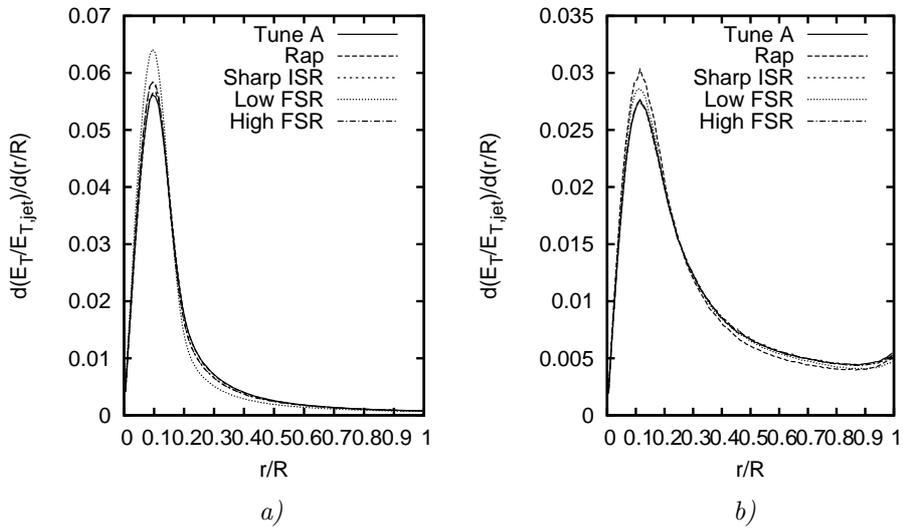


Figure 19: The *a)* hardest and *b)* third hardest jet E_{\perp} profile for 1.96 TeV $p\bar{p}$ events with $p_{\perp\text{hard}} > 100$ GeV.

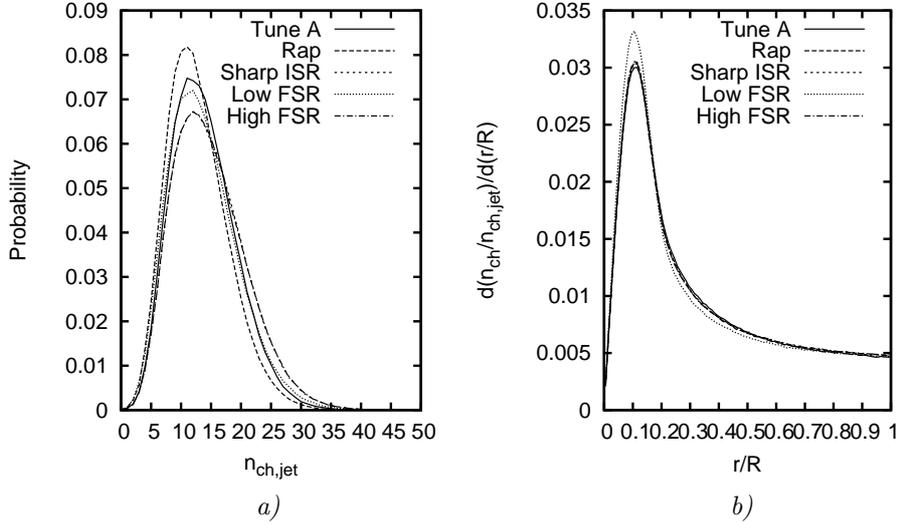


Figure 20: The *a*) charged multiplicity distribution and *b*) charged particle profile of the hardest jet in 1.96 TeV $p\bar{p}$ events with $p_{\perp\text{hard}} > 100$ GeV.

to be quite similar. The problem therefore remains that this primordial k_{\perp} is larger than can physically be well motivated based on purely nonperturbative physics. We observe that, among the new models, the Sharp ISR could have been combined with a smaller primordial k_{\perp} since its peak is shifted towards too large $p_{\perp Z}$, while the High FSR and Low FSR (which here only differ by their $p_{\perp 0}$ values) could have used an even larger primordial k_{\perp} . In part, this makes sense: with ISR being turned off at larger p_{\perp} values in the latter models, it is then also easier to motivate a larger primordial k_{\perp} .

The complete comparison of algorithms is rather complicated, however. The primordial k_{\perp} that reaches the hard interaction is diluted by the ISR activity, and so scales down like the ratio of the x value of the incoming parton at the hard interaction to that of the initiator, $z_{\text{tot}} = x_{\text{in}}/x_{\text{init}}$. This ratio is approximately the same for the old ISR shower ($\langle z_{\text{tot}} \rangle \approx 0.59$) and Sharp ISR ($\langle z_{\text{tot}} \rangle \approx 0.62$), indicating that the fewer ISR branchings and smaller z per branching in the new algorithm rather well cancel. The smooth turnoff of High and Low FSR gives less branchings ($\langle z_{\text{tot}} \rangle \approx 0.75$) and thus more primordial k_{\perp} survives in these scenarios.

In summary, the new MI+ISR scheme gives an improved description of Z^0 production, but does not remove the need for an uncomfortably large primordial k_{\perp} .

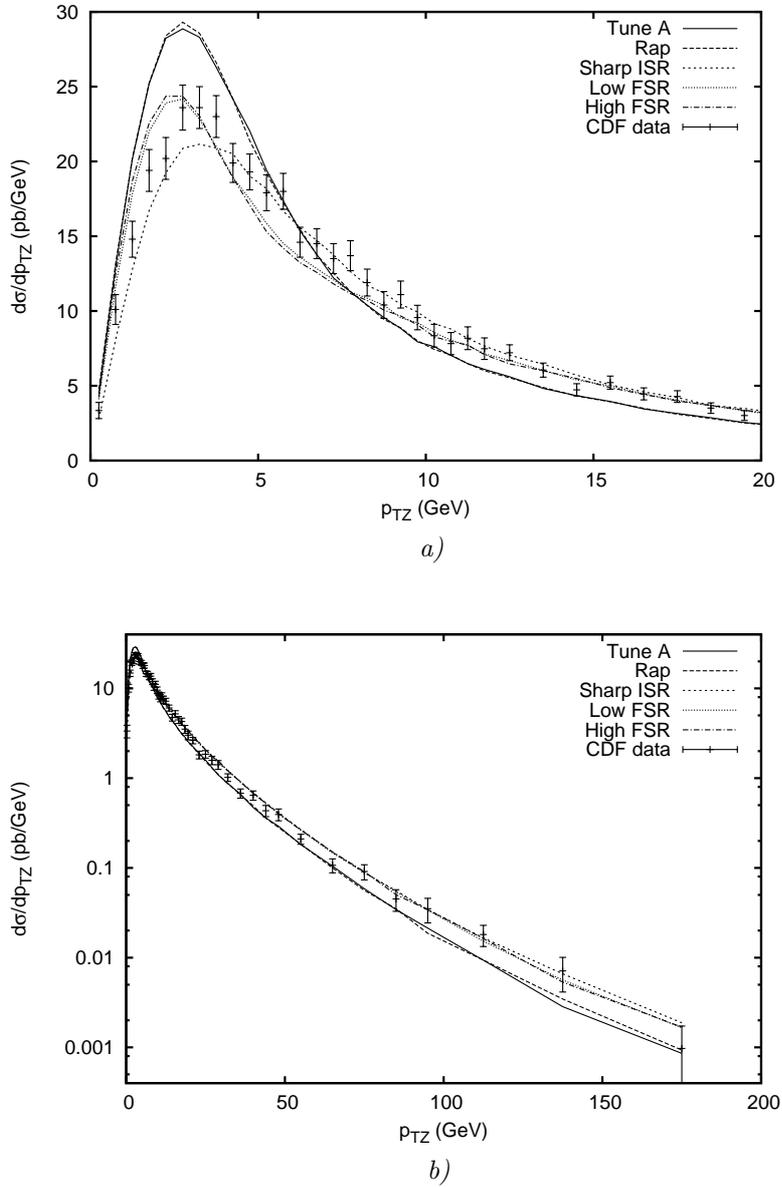


Figure 21: The p_{\perp} spectrum at *a*) low and *b*) all p_{\perp} for dilepton pairs in the 66–116 GeV mass range, γ^*/Z^0 simulation compared with CDF corrected data at 1.8 TeV [35].

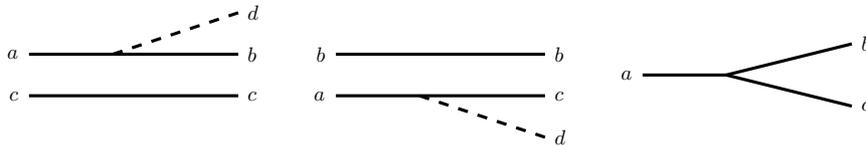


Figure 22: Illustration of the three terms in the two-parton density evolution, eq. (46).

5 Outlook

In this article we have considered the consequences of interleaved multiple interactions and initial-state radiation, and paved the way for interleaving also final-state radiation in this framework, but that does not exhaust the list of perturbative processes in ‘normal’ hadronic events. One further possibility is that a parton from one of the incoming hadrons scatters twice, against two different partons from the other hadron, rescattering or ‘ $3 \rightarrow 3$ ’. Another possibility is that two partons participating in two separate hard scatterings may turn out to have a common ancestor when the backwards evolution traces the prehistory to the hard interactions, *joined interactions* (JI).

The $3 \rightarrow 3$ processes have been considered in the literature [36], with the conclusion that they should be less important than multiple $2 \rightarrow 2$ processes, except possibly at large p_{\perp} values, where QCD radiation anyway is expected to be the dominant source of multijet events. The reason is that one $3 \rightarrow 3$ scattering and two $2 \rightarrow 2$ ones have similar parton-level cross sections, but the latter wins by involving one parton density more. Nevertheless, at some point, there ought to be a more detailed modelling, in order better to quantify effects.

The joined interactions are well-known in the context of the evolution of multiparton densities [37], but have not been applied to a multiple interactions framework. We will therefore here carry out a first study, to quantify roughly how common JI are and how much activity they contribute with. A full implementation of the complete kinematics, intertwining MI, ISR and JI all possible ways, is a major undertaking, worth the effort only if the expected effects are non-negligible. Given the many uncertainties in all the other processes at play, one would otherwise expect that the general tuning of MI/ISR/FSR/... to data would hide the effects of JI, as well as of $3 \rightarrow 3$ processes.

5.1 Joined interactions: theory

Just like the starting point for a discussion of ISR is the DGLAP evolution equations for the single-parton densities, the starting point for JI is the evolution equations for the two-parton densities. Relevant forwards-evolution formulae are available in the literature in integrated form [37]. Here, however, we will

choose a differential form, that can then be applied to our backwards evolution framework. To this end, define the two-parton density $f_{bc}^{(2)}(x_b, x_c, Q^2)$ as the probability to have a parton b at energy fraction x_b and a parton c at energy fraction x_c when the proton is probed at a scale Q^2 . The evolution equation for this distribution is

$$\begin{aligned} df_{bc}^{(2)}(x_b, x_c, Q^2) = & \frac{dQ^2}{Q^2} \frac{\alpha_s}{2\pi} \iint dx_a dz \left\{ \right. \\ & f_{ac}^{(2)}(x_a, x_c, Q^2) P_{a \rightarrow bd}(z) \delta(x_b - zx_a) \\ & + f_{ba}^{(2)}(x_b, x_a, Q^2) P_{a \rightarrow cd}(z) \delta(x_c - zx_a) \\ & \left. + f_a(x_a, Q^2) P_{a \rightarrow bc}(z) \delta(x_b - zx_a) \delta(x_c - (1-z)x_a) \right\} \end{aligned} \quad (46)$$

As usual, we assume implicit summation over the allowed flavour combinations; thus the last term is absent when there is no suitable mother a for a given set of b and c . An illustration of the three terms is given in Fig. 22. The first two are the standard ones, where b and c evolve independently, up to flavour and momentum conservation constraints, and are already taken into account in the ISR framework. It is the last term that describes the new possibility of two evolution chains having a common ancestry.

Carrying out the δ integrations, which imply that $x_a = x_b + x_c$ and $z = x_b/(x_b + x_c)$, the probability for the unresolution of b and c into a when Q^2 is decreased (cf. the step from eq. (6) to eq. (7)) can be rewritten as

$$\begin{aligned} d\mathcal{P}_{bc}(x_b, x_c, Q^2) &= \left| \frac{df_{bc}^{(2)}(x_b, x_c, Q^2)}{f_{bc}^{(2)}(x_b, x_c, Q^2)} \right| \\ &= \left| \frac{dQ^2}{Q^2} \right| \frac{\alpha_s}{2\pi} \frac{f_a(x_a, Q^2)}{f_{bc}^{(2)}(x_b, x_c, Q^2)} \frac{1}{x_b + x_c} P_{a \rightarrow bc}(z) \\ &= \left| \frac{dQ^2}{Q^2} \right| \frac{\alpha_s}{2\pi} \frac{x_a f_a(x_a, Q^2)}{x_b x_c f_{bc}^{(2)}(x_b, x_c, Q^2)} z(1-z) P_{a \rightarrow bc}(z) \\ &\simeq \left| \frac{dQ^2}{Q^2} \right| \frac{\alpha_s}{2\pi} \frac{x_a f_a(x_a, Q^2)}{x_b f_b(x_b, Q^2) x_c f_c(x_c, Q^2)} z(1-z) P_{a \rightarrow bc}(z) . \end{aligned} \quad (47)$$

In the last step we have introduced the approximation $f_{bc}^{(2)}(x_b, x_c, Q^2) \simeq f_b(x_b, Q^2) f_c(x_c, Q^2)$ to put the equation in terms of more familiar quantities. Just like for the other processes considered, a form factor is given by integration over the relevant Q^2 range and exponentiation.

The strategy now is clear. Previously we have introduced a scheme wherein events are evolved downwards in p_\perp . At each step a new trial multiple interaction competes against trial ISR branchings on the existing interactions, and

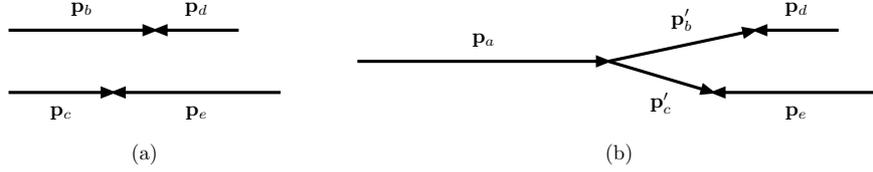


Figure 23: Kinematics of the $b + d$ and $c + e$ colliding systems (a) before and (b) after the $a \rightarrow b + c$ branching is reconstructed.

the one with largest p_{\perp} ‘wins’. Now a third option is added, competing with the first two in the same way, i.e. eq. (45) is extended to

$$\begin{aligned} \frac{d\mathcal{P}}{dp_{\perp}} &= \left(\frac{d\mathcal{P}_{\text{MI}}}{dp_{\perp}} + \sum \frac{d\mathcal{P}_{\text{ISR}}}{dp_{\perp}} + \sum \frac{d\mathcal{P}_{\text{JI}}}{dp_{\perp}} \right) \times \\ &\times \exp \left(- \int_{p_{\perp}}^{p_{\perp}^{i-1}} \left(\frac{d\mathcal{P}_{\text{MI}}}{dp'_{\perp}} + \sum \frac{d\mathcal{P}_{\text{ISR}}}{dp'_{\perp}} + \sum \frac{d\mathcal{P}_{\text{JI}}}{dp'_{\perp}} \right) dp'_{\perp} \right). \end{aligned} \quad (48)$$

The JI sum runs over all pairs of initiator partons with allowable flavour combinations, separately for the two incoming hadrons. A gluon line can always be joined with a quark or another gluon one, and a sea quark and its companion can be joined into a gluon. For each of these possibilities, $d\mathcal{P}_{\text{JI}} \exp(-\int d\mathcal{P}_{\text{JI}})$ can be used to do a backwards evolution from the $p_{\perp \text{max}} = p_{\perp i-1}$ scale given by the previous step. If such a trial joining occurs at a larger p_{\perp} scale than any of the other trial possibilities, then it is allowed to occur. Also the regularization procedure at small p_{\perp} values is the same as for MI and ISR.

The parton densities we will use are defined in the same spirit as previously discussed, e.g. $f_b(x_b, p_{\perp}^2)$ and $f_c(x_c, p_{\perp}^2)$ are squeezed into ranges $x \in [0, X]$, where X is reduced from unity by the momentum carried away by all but the own interaction, and for $f_a(x_a, p_{\perp}^2)$ by all but the b and c interactions. Note that companion distributions are normalized to unity. Therefore, for heavy quarks, the branching probability $g \rightarrow Q\bar{Q}$ goes like $1/\ln(p_{\perp}^2/m_Q^2)$ for $p_{\perp}^2 \rightarrow m_Q^2$, as it should, rather than like $1/\ln^2(p_{\perp}^2/m_Q^2)$, which would have been obtained if f_Q and $f_{\bar{Q}}$ independently were assumed to vanish in this limit.

Unfortunately the kinematics reconstruction offers a complication. Consider a system with recoilers d and e to b and c , respectively, as depicted in Fig. 23. Use a prime to denote the modified b and c four-momenta after the $a \rightarrow b + c$ branching has been constructed, while d and e should be unchanged. From $x_a = x_b + x_c$ it then follows that $p_a = p_b + p_c = p'_b + p'_c$, and from the z definition that

$$(p'_b + p_d)^2 = (p_b + p_d)^2 = z(p_a + p_d)^2, \quad (49)$$

$$(p'_c + p_e)^2 = (p_c + p_e)^2 = (1-z)(p_a + p_e)^2. \quad (50)$$

Further, p'_b and p'_c should have opposite and compensating transverse momenta given by the p_\perp scale above, and spacelike virtualities to be determined. Now, it turns out that these requirements overconstrain the system. The basic problem is illustrated by eq. (27): the spacelike parton needs to pick up a larger p_\parallel component than its z share, in order to retain the invariant mass with the recoiler when the p_\perp is introduced. So, if both daughters should be spacelike, not both of them can pick up more p_\parallel than E , given that a is massless. (A solution where one of p'_b and p'_c is timelike sometimes exists, but not always, and is anyway rather contrived.)

We see two alternative ways out of this dilemma.

- Retain the $x_a = x_b + x_c$ expression, at the expense of not giving any p_\perp or virtualities in the branching, i.e. $p'_b = p_b$ and $p'_c = p_c$. Then p_\perp only plays the role of a formal evolution parameter, denoting the scale above which b and c may radiate and interact separately.
- Insist on having a p_\perp kick in the branching. Then a sensible (but not unique) choice is to put $p'^2_b = p'^2_c = -p_\perp^2$, such that both have $m_\perp = 0$ and thus $p_\parallel = E$. These energies must now be scaled up somewhat, to $E'_b = (1 + p_\perp^2/m_{bd}^2)E_b$ and $E'_c = (1 + p_\perp^2/m_{ce}^2)E_c$, for the invariant masses with the recoiler to be preserved, and therefore

$$x_a = \left(1 + \frac{p_\perp^2}{m_{bd}^2}\right) x_b + \left(1 + \frac{p_\perp^2}{m_{ce}^2}\right) x_c . \quad (51)$$

It is then this x_a that should be used in parton densities, to ensure that the probability of a joining is suppressed near the kinematical limit.

Given that no joinings are possible until after (at least) two interactions have been generated, and that the rate increases roughly quadratically with the number of interactions, this physics mechanism becomes more important at smaller p_\perp values. Therefore we do not expect the above two extremes to differ that significantly for practical applications.

5.2 Joined interactions: results

Although an algorithm implementing the full kinematics for joined interactions has not yet been constructed, it is still possible to gauge the order of magnitude of the effects such joinings could have. We do this by formally performing the backwards evolution according to eq. (48), i.e. including the joining term eq. (47) in competition with the ordinary ISR and MI terms, without actually letting the generated joinings occur physically. Thereby we still obtain an estimate for how often and at which p_\perp values joinings would occur.

Since we do not perform the joinings physically, the backwards evolution could in principle attempt joinings involving the *same* initial state shower chain more than once. Such joinings are of course rejected; only the first joining involving a particular chain is kept track of.

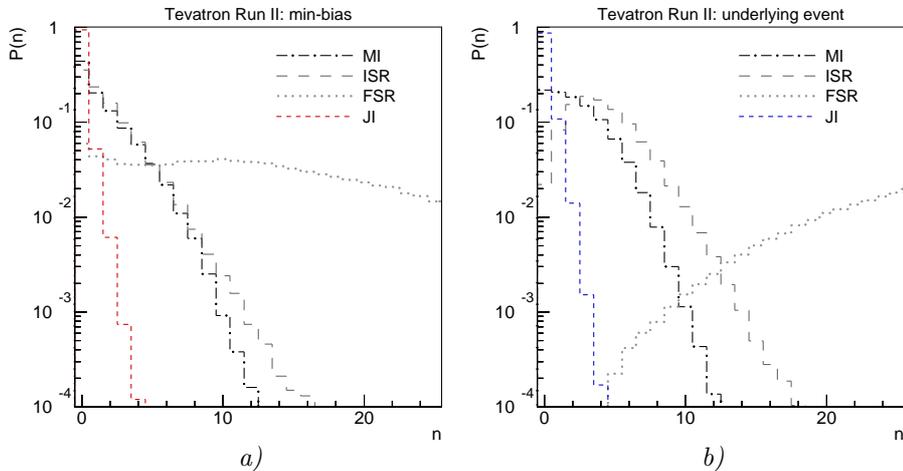


Figure 24: Probability distributions of MI, ISR, FSR, and JI in *a)* min-bias events and *b)* events with $p_{\perp\text{hard}} > 100$ GeV, for 1.96 TeV $p\bar{p}$ events. Note that the MI distribution does *not* include the hardest scattering.

Taking the High FSR model in Table 1 as a fair representative of the evolution in the new framework, we show the number distributions of multiple interactions excluding the first (MI), ISR branchings, FSR branchings, and trial joinings (JI) in 1.96 TeV $p\bar{p}$ minimum-bias events, Fig. 24a, and for events where the p_{\perp} of the hard interaction is above 100 GeV, Fig. 24b. Below, we refer to the former as the “min-bias” sample and to the latter as the “UE” (underlying event) sample. FSR is shown mainly for reference here, the important graphs being the ones illustrating the evolution in the initial state: MI, ISR, and JI. One clearly observes that joinings are much less frequent than the other types of evolution steps, averaging at roughly one joining per 15 events for the min-bias sample and one per 7 events for the UE sample. Thus, even when relatively hard physics is involved, shower joinings do not appear to take a very prominent role in the evolution.

To complement the number distributions, Fig. 25 shows *where* the evolution steps occur in $p_{\perp\text{evol}}$. As expected, the joinings occur at comparatively low values of $p_{\perp\text{evol}}$. Also notice that both the ISR, MI, and JI distributions exhibit a turnover around $p_{\perp 0}$, characteristic of the smooth regularization used in the High FSR model.

Finally, Fig. 26 shows the total p_{\perp} sum of MI, ISR, FSR, and JI activity, respectively. That is, for each interaction, branching, or joining, a scalar p_{\perp} is defined, which is added to a cumulative sum. For MI and JI this p_{\perp} is defined with respect to the beam axis, while the ISR and FSR p_{\perp} is defined with respect to the branching parton, which for ISR is roughly along the beam

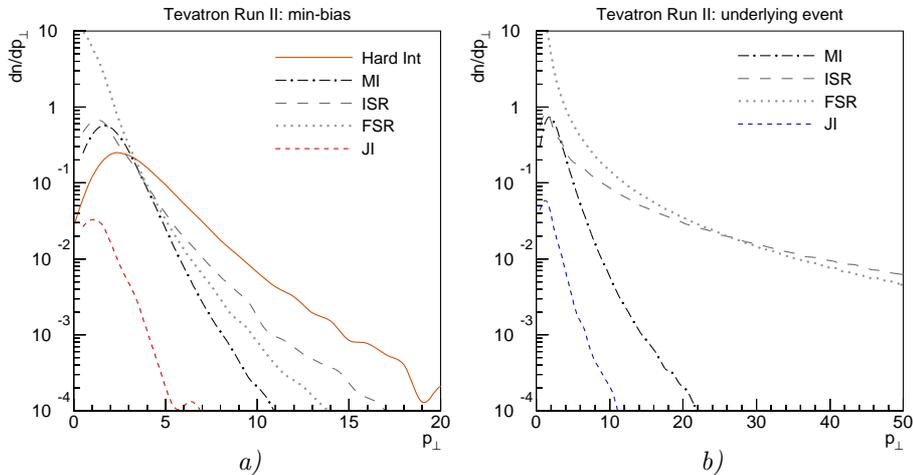


Figure 25: $p_{\perp, \text{evol}}$ distributions showing the scale at which multiple interactions (MI), ISR branchings, FSR branchings, and joined interactions (JI) occur, in 1.96 TeV $p\bar{p}$ min-bias events. *a)* Minimum-bias events, with the p_{\perp} scale of the hardest interaction shown for reference (solid line). *b)* events with $p_{\perp, \text{hard}} > 100$ GeV. Note that the p_{\perp} axis goes out to 20 GeV in *a)* and to 50 GeV in *b)*.

direction, but for FSR normally not. Therefore FSR mainly broadens jets, i.e. redistributes the existing \cancel{E}_T , whereas the other mechanisms increase the total \cancel{E}_T of the event. Again, the cumulative effect of joinings is small, with only about 1% of the min-bias sample and 5% of the UE one exhibiting more than 2 GeV of total p_{\perp} from joinings.

6 Conclusions

It would seem natural to consider the evolution of a high-energy event in the normal time order. In such a framework, the incoming hadrons are evolved from a simple partonic configuration at a low Q_0 scale, up through a number of short-lived fluctuations at different virtuality scales. At the moment of collision, the two sets of partons may undergo several independent interactions. The scattered partons can radiate in the final state, while fluctuations inside which no scatterings occurred may recombine. Finally the set of outgoing low-virtuality partons hadronize.

Such an approach has the advantage that it inherently provides multiparton distributions, and thereby automatically contains correlations between interactions, including what is here called joined interactions. It does not remove the need to consider possible scatterings in some order of hardness, however:

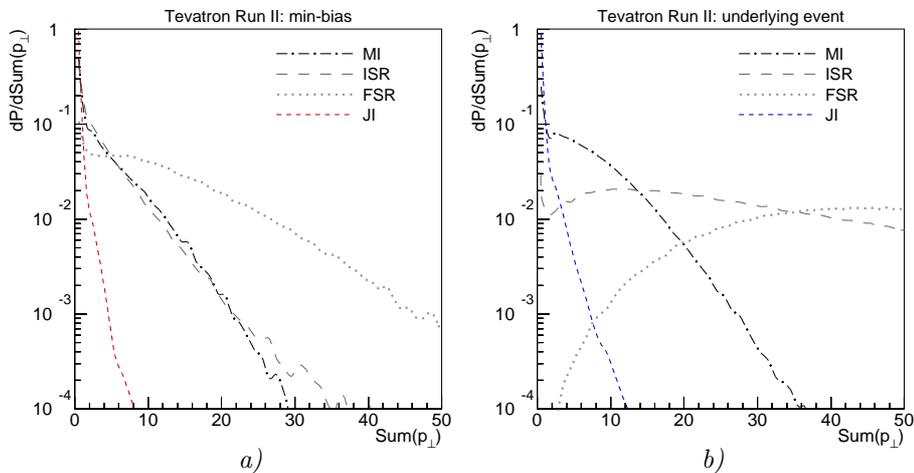


Figure 26: $\sum p_{\perp}$ distributions for MI, ISR, FSR and JI, in 1.96 TeV $p\bar{p}$ *a)* minimum-bias events and *b)* events with $p_{\perp\text{hard}} > 100$ GeV. Note that the MI distribution does not include the hardest scattering.

the partons of a high-virtuality fluctuation may either interact individually or collectively, the latter as the unresolved mother parton at a lower resolution scale, and the former should preempt the latter.

There is also the well-known problem that it is difficult to generate rare processes, since there is no straightforward way to preselect the forwards evolution to give the desired configuration. The nightmare example is the production of a narrow Higgs state, where the incoming partons must match very precisely in invariant mass for a reaction to be possible. More generally, efficiency suffers from the need to consider a wealth of virtual fluctuations that in the end lead to nothing. The assignment of individual virtualities and transverse momenta to partons in such fluctuations is also not unique, and does affect the kinematics reconstruction. And, of course, the whole plethora of coherence effects need to be considered.

The alternative is to start with the hardest interaction, and then ‘work outwards’ to longer timescales both in the past and future, i.e. to (re)construct less hard steps in the evolution of the event. This makes the preselection of desired events straightforward, and in general implies that the most striking aspects of the event are considered ‘up front’. The price is a tougher task of reconstructing the soft associated activity in the initial state, while final-state radiation and hadronization offer about the same challenges in the two scenarios.

In a set of articles we have begun the task of providing an improved description of events along this latter philosophy. The first step [3] was to develop

a model for the hadronization of junction topologies, thereby allowing complicated beam remnants. The second step [1] was to develop a framework for correlated parton densities, in flavour, colour and transverse and longitudinal momentum, thereby allowing initial-state radiation to be considered in full. In this, the third step, we have interleaved multiple interactions and initial-state radiation in a common transverse-momentum-ordered sequence, with a common dampening procedure at low transverse momenta to handle destructive interference in this region. Our lack of explicitly implemented joined interactions has been shown not to be a major shortcoming, since such joinings are reasonably rare. That is, taken together, we now have a framework that should provide a complete description of all aspects that could have been covered by a forwards-in-time evolution scenario, along with the traditional advantages of the backwards-evolution approach. In addition, the new framework makes use of new algorithms for p_{\perp} -ordered evolution in initial- and final-state radiation, which should further improve the quality of the description.

It may then be somewhat disappointing that we here have used PYTHIA Tune A [34] as a reference, well knowing that Tune A is able to describe a host of jet and minimum-bias data at the Tevatron, in spite of it being based on a much more primitive approach [2]. The hope, of course, is that our new approach will be able to explain — and predict — much more data than Tune A can. For sure we know of many aspects of the old framework that are unreasonable, but that either have not been probed or that may have been fixed up by a contrived choice of tuned parameters. Ultimately this is for experimentalists to tell, as tests become increasingly more sophisticated. Certainly, one should not expect the advantages of the new model to become apparent unless a similar effort is mounted as went into producing Tune A in the first place.

There are also a few issues still hanging over us, awaiting a ‘fourth step’. One is the implementation of joined interactions and $3 \rightarrow 3$ rescattering processes, to see what their real impact is, whether negligible or not. But the main one we believe to be the interleaving of final-state radiation with multiple interactions and initial-state radiation. On the one hand, such an interleaving may not be required, since the competition between FSR and MI+ISR is less direct than that between MI and ISR: an FSR emission at a high p_{\perp} scale does not affect the probability for MI or ISR activity at lower p_{\perp} values. On the other hand, there would then also not seem to be any disadvantage to having a commonly ordered p_{\perp} sequence of MI+ISR+FSR, and such an ordering would come in handy for a consistent interfacing to higher-order matrix elements. Furthermore, a p_{\perp} -ordered FSR algorithm is available, well matched to the p_{\perp} -ordering of MI and ISR.

There is, however, one major open question related to FSR interleaving: which parton takes the momentum recoil when a FSR branching pushes a parton off the mass shell? The problem is not so much the momentum transfer itself, but that the size of the radiating dipole sets the maximum scale for

allowed emissions. We have in this article illustrated how such a choice can affect e.g. the jet multiplicity and jet profiles. The crucial distribution is the $\langle p_{\perp} \rangle (n_{\text{ch}})$ one, however. In order to provide a reasonable description of the experimental data, we are forced to arrange colours in the final state to have a smaller string length than colour correlations in the initial state alone would suggest. This problem has ‘always’ been there [2], and is accentuated in Tune A, where as much as 90% of the partons added by multiple interactions are connected so as to minimize the string length. The hope that an improved treatment of other aspects would remove the need for a special string-length minimization mechanism has so far failed to materialize. We therefore need to understand better how the colour flow is set, and how this influences the evolution of an event, especially the FSR activity.

The related fields of minimum bias physics and underlying events thus are further explored but not solved with this article, and likely not with the next one either. This should come as no surprise: in the world of hadronic physics, there are few simple answers. Everything that is not explicitly forbidden is bound to happen, and often at a significant rate. To reflect reality, the theoretical picture therefore has to become more and more complex, as one consideration after the next is pulled into the game. However, if the journey is interesting and educational, why despair that the end station is not yet reached?

Acknowledgements

We are grateful to G. Rudolph for having tuned our new p_{\perp} -ordered final-state shower algorithm to LEP data. Communications with Tevatron and LHC physicists, especially R.D. Field, J. Huston and A. Moraes, have spurred us on.

References

- [1] T. Sjöstrand and P.Z. Skands, JHEP **03** (2004) 053
- [2] T. Sjöstrand and M. van Zijl, Phys. Rev. **D36** (1987) 2019
- [3] T. Sjöstrand and P.Z. Skands, Nucl. Phys. **B659** (2003) 243
- [4] T. Sjöstrand, Phys. Lett. **157B** (1985) 321;
M. Bengtsson, T. Sjöstrand and M. van Zijl, Z. Phys. **C32** (1986) 67
- [5] T. Sjöstrand, LU TP 04–05 [hep-ph/0401061], in the QCD/SM Working Group Summary Report of the Physics at TeV Colliders Workshop, Les Houches, France, 26 May – 6 Jun 2003 [hep-ph/0403100]
- [6] T. Sjöstrand, P. Edén, C. Friberg, L. Lönnblad, G. Miu, S. Mrenna and E. Norrbin, Computer Phys. Commun. **135** (2001) 238;

- T. Sjöstrand, L. Lönnblad, S. Mrenna and P. Skands, LU TP 03-38 [hep-ph/0308153].
- [7] V.N. Gribov and L.N. Lipatov, Sov. J. Nucl. Phys. **15** (1972) 438, *ibid.* 75;
G. Altarelli and G. Parisi, Nucl. Phys. **B126** (1977) 298;
Yu.L. Dokshitzer, Sov. J. Phys. JETP **46** (1977) 641
- [8] B. Andersson, G. Gustafson, G. Ingelman and T. Sjöstrand, Phys. Rep. **97** (1983) 31;
T. Sjöstrand, Nucl. Phys. **B248** (1984) 469;
B. Andersson, 'The Lund Model' (Cambridge University Press, 1998)
- [9] V.V. Sudakov, Zh.E.T.F. **30** (1956) 87 (Sov. Phys. J.E.T.P. **30** (1956) 65)
- [10] M. Bengtsson and T. Sjöstrand, Nucl. Phys. **B289** (1987) 810
- [11] E. Norrbin and T. Sjöstrand, Nucl. Phys. **B603** (2001) 297
- [12] G. Marchesini and B.R. Webber, Nucl. Phys. **B238** (1984) 1;
G. Corcella, I.G. Knowles, G. Marchesini, S. Moretti, K. Odagiri, P. Richardson, M.H. Seymour and B.R. Webber, JHEP **01** (2001) 010, hep-ph/0210213
- [13] G. Gustafson, Phys. Lett. **B175** (1986) 453
- [14] G. Gustafson and U. Pettersson, Nucl. Phys. **B306** (1988) 746
- [15] L. Lönnblad, Computer Physics Commun. **71** (1992) 15
- [16] A.H. Mueller, Phys. Lett. **104B** (1981) 161;
B.I. Ermolaev, V.S. Fadin, JETP Lett. **33** (1981) 269
- [17] M.H. Seymour, Computer Physics Commun. **90** (1995) 95
- [18] G. Gustafson, private communication
- [19] S. Catani, F. Krauss, R. Kuhn and B.R. Webber, JHEP **11** (2001) 063;
L. Lönnblad, JHEP **05** (2002) 046
- [20] P. Skands and T. Sjöstrand, LU TP 03-45 [hep-ph/0310315], in the proceedings of International Europhysics Conference on High-Energy Physics (HEP 2003), Aachen, Germany, 17-23 Jul 2003, European Physical Journal **C** Direct, DOI: 10.1140/epjcd/s2003-03-520-7;
T. Sjöstrand and P. Skands, LU TP 04-04 [hep-ph/0401060], in The QCD/SM working group: Summary report, 3rd Les Houches Workshop: Physics at TeV Colliders, Les Houches, France, 26 May - 6 Jun 2003, M. Dobbs *et al.* [hep-ph/0403100]

- [21] OPAL Collaboration, M.Z. Akrawy et al., *Z. Phys.* **C47** (1990) 505;
L3 Collaboration, B. Adeva et al., *Z. Phys.* **C55** (1992) 39;
I.G. Knowles et al., in ‘Physics at LEP2’, eds. G. Altarelli, T. Sjöstrand and F. Zwirner, CERN 96–01 (Geneva, 1996), Vol. 2, p. 103;
DELPHI Collaboration, P. Abreu et al., *Z. Phys.* **C73** (1996) 11;
ALEPH Collaboration, R. Barate et al., *Phys. Rep.* **294** (1998) 1
- [22] B. Andersson et al., *Eur. Phys. J.* **C25** (2002) 77;
J. Andersen et al., *Eur. Phys. J.* **C35** (2004) 67
- [23] B. Andersson, G. Gustafson, L. Lönnblad and U. Pettersson, *C. Phys.* **C43** (1989) 625
- [24] B. Andersson, G. Gustafson and J. Samuelsson, *Nucl. Phys.* **B467** (1996) 443;
B. Andersson, G. Gustafson and H. Kharraziha, *Phys. Rev.* **D57** (1998) 5543;
H. Kharraziha and L. Lönnblad, *JHEP* **03** (1998) 006
- [25] M. Ciafaloni, *Nucl. Phys.* **B296** (1987) 249;
S. Catani, F. Fiorani and G. Marchesini, *Nucl. Phys.* **B336** (1990) 18
- [26] T. Sjöstrand, *Computer Physics Commun.* **28** (1983) 229
- [27] S. Catani, Yu. L. Dokshitzer, M. Olsson, G. Turnock and B.R. Webber, *Phys. Lett.* **B269** (1991) 432
- [28] D. Amati, A. Bassetto, M. Ciafaloni, G. Marchesini and G. Veneziano, *Nucl. Phys.* **B173** (1980) 429;
G. Curci, W. Furmanski and R. Petronzio, *Nucl. Phys.* **B175** (1980) 27
- [29] B.R. Webber, *Ann. Rev. Nucl. Part. Sci.* **36** (1986) 253
- [30] G. Rudolph, private communication
- [31] G. Miu and T. Sjöstrand, *Phys. Lett.* **B449** (1999) 313
- [32] E. Thomé, master’s thesis, Lund University, LU TP 04–01 [hep-ph/0401121];
J. Huston, I. Puljak, T. Sjöstrand, E. Thomé, LU TP 04–07 [hep-ph/0401145], in the QCD/SM Working Group Summary Report of the Physics at TeV Colliders Workshop, Les Houches, France, 26 May – 6 Jun 2003 [hep-ph/0403100]
- [33] A. Breakstone et al., *Phys. Lett.* **132B** (1983) 463;
UA1 Collaboration, C. Albajar et al., *Nucl. Phys.* **B335** (1990) 261;
E735 Collaboration, T. Alexopoulos et al., *Phys. Lett.* **B336** (1994) 599;
CDF Collaboration, D. Acosta et al., *Phys. Rev.* **D65** (2002) 072005

- [34] R.D. Field, presentations at the ‘Matrix Element and Monte Carlo Tuning Workshop’, Fermilab, 4 October 2002 and 29–30 April 2003, talks available from webpage <http://cepa.fnal.gov/CPD/MCTuning/>, and further recent talks available from <http://www.phys.ufl.edu/~rfield/cdf/>
- [35] CDF Collaboration, T. Affolder et al., Phys. Rev. Lett. **84** (2000) 845
- [36] N. Paver and D. Treleani, Nuovo Cimento **70A** (1982) 215, Phys. Lett. **146B** (1984) 252, Z. Phys. **C28** (1985) 187
- [37] K. Konishi, A. Ukawa and G. Veneziano, Nucl. Phys. **B157** (1979) 45;
R. Kirschner, Phys. Lett. **84B** (1979) 266;
V.P. Shelest, A.M. Snigirev and G.M. Zinovjev, Phys. Lett. **113B** (1982) 325;
A.M. Snigirev, Phys. Rev. **D68** (2003) 114012;
V.L. Korotkikh and A.M. Snigirev, hep-ph/0404155