Introduction to the LHC

Michelangelo M. Mangano – *TH Division, CERN. The scientific interests include: Particle Physics, Physics at LHC, QCD.*

Abstract: This is a condensed summary of a series of lectures on the Large Hadron Collider (LHC). We briefly review its physics goals, the accelerator itself, and the basic principles underlying the description of hadronic collisions.

Introduction to the LHC

Michelangelo M. Mangano[a]

[a] *PH Department, TH Unit, CERN, CH-1211 Geneva 23, Switzerland*

Abstract

This is a condensed summary of a series of lectures on the Large Hadron Collider (LHC). We briefly review its physics goals, the accelerator itself, and the basic principles underlying the description of hadronic collisions.

1 The main goals of the LHC programme

The Standard Model (SM), whose complete formulation dates back to the early 70's, has been shown over the past 30 years to accurately describe all properties of the interactions among fundamental particles, namely quarks, leptons and the gauge bosons transmitting the electroweak and strong forces. Its internal consistency, nevertheless, relies on a mechanism to break the symmetry between electromagnetic and weak interactions, the so-called electroweak symmetry breaking (EWSB). To this end, the SM assumes the existence of a scalar field with a weak charge, the Higgs H, whose potential energy is minimized with a non-vanishing value of its matrix element on the vacuum state, $\langle H \rangle = v \neq 0$. This leads to spontaneous symmetry breaking. The measured strength of the weak interactions and the mass of their carriers, the $W^{\pm}$ gauge bosons, fix the value of $v \sim 247$ GeV, thus setting the natural mass scale for weak phenomena. The fluctuations of the Higgs field around its vacuum state give rise to a particle, the Higgs boson, whose mass m_H is a free parameter of the model. Direct searches at the LEP e^+e^- collider have established a lower limit $m_H \gtrsim 114$ GeV, important constraints in the mass range around 170 GeV have been recently achieved by the Tevatron experiments, and theoretical analyses of the consistency of the model set an upper limit, around 800 GeV.

The Higgs boson predicted by the SM, therefore, has not yet been discovered and the EWSB mechanism has yet to be verified experimentally. The main priority of particle physics today is to confirm this mechanism ... or to discover what replaces it! There are in fact good reasons for theorists to suspect that physics *beyond* the SM (BSM), and beyond the basic Higgs mechanism outlined above, should play a key role in the dynamics of EWSB [1]. To start with, one would like to have a framework within which the smallness of the weak scale $v \sim 247$ GeV, relative to the Planck scale

$M_{Pl} \sim 10^{19}$ GeV, is the natural result of dynamics, as opposed to a random accident. Furthermore, there are unequivocal experimental indications that BSM phenomena are required to explain what is observed in the universe: the SM can explain neither the existence of dark matter, nor the ratio of baryons and radiation present in the universe. In addition, although neutrino masses could be incorporated with a minimal and trivial adjustment of the SM spectrum, the most compelling explanations of how neutrinos acquire such a small mass rely on the existence of new phenomena at scales of the order of the Grand Unification, $M_{GUT} \sim 10^{15}$ GeV. Furthermore, there are several questions that within the SM cannot be addressed, but that could acquire a dynamical content in a broader framework. As an example, consider the issue of what is the origin of the three quark and lepton generations and of the diverse mass patterns between and within them. The various BSM proposals anticipate new phenomena to appear at the TeV mass scale: supersymmetry, which implies the existence of a new boson for each fermionic particle of the SM, and of a new fermion for each known boson; new forces, mediated by heavy gauge bosons, possibly restoring at high energy the chiral left-right asymmetry of the low-energy world; new strong interactions and new fermions, responsible for the dynamical generation of EWSB; compactified space dimensions, leading to the existence of an infinite Kaluza-Klein spectrum of particles corresponding to each known one, with ever increasing and linearly-spaced masses; composite structures within what are considered as fundamental, elementary particles; and more, with a continuous emergence of new ideas and proposals to embed the SM into a more complete theory.

The Large Hadron Collider (LHC) will be the first accelerator operating at energies high enough to explore a large fraction of these proposals. Its design, and that of its two largest experiments, ATLAS [2] and CMS [3], has been tuned to enable the full exploration of the Higgs mass range, searching for a broad variety of the Higgs production and decay processes predicted by the SM and beyond. This short summary of my lectures will start by recalling the basic principles underlying our understanding of hadronic collisions, the basic element to predict rates and properties of the production of new particles such as the Higgs. I will then give a brief account of the accelerator parameters, and of the expected timeline for discoveries.

There is a vast literature behind the topics I shall just touch upon. But for more details, at an accessible level, I refer you to some basic textbook or review articles, such as [4] for the QCD aspects, and [5] for an introduction to the SM.

2 The proton structure and hadronic collisions

The understanding of the structure of the proton at short distances is one of the key ingredients to predict cross sections for processes involving hadrons in the initial state. All processes in hadronic collisions, even those intrinsically of electroweak

nature such as the production of W/Z bosons or photons, are in fact induced by the quarks and gluons contained inside the hadron. In this section I shall introduce some important concepts, such as the notion of partonic densities of the proton, and of parton evolution.

We shall limit ourselves to processes where a proton–(anti)proton pair collides at large centre-of-mass energy ($\sqrt{S}$, typically larger than several hundred GeV) and undergoes a very inelastic interaction, with momentum transfers between the participants in excess of several GeV. The outcome of this hard interaction could be the simple scattering at large angle of some of the hadron's elementary constituents, their annihilation into new massive resonances, or a combination of the two. In all cases the final state consists of a large multiplicity of particles, associated to the evolution of the fragments of the initial hadrons, as well as of the new states produced. As discussed below, the fundamental physical concept that makes the theoretical description of these phenomena possible is 'factorization', namely the ability to isolate separate independent phases of the overall collision. These phases are dominated by different dynamics, and the most appropriate techniques can be applied to describe each of them separately. In particular, factorization allows one to decouple the complexity of the proton structure and of the final-state hadron formation from the elementary nature of the perturbative hard interaction among the partonic constituents.

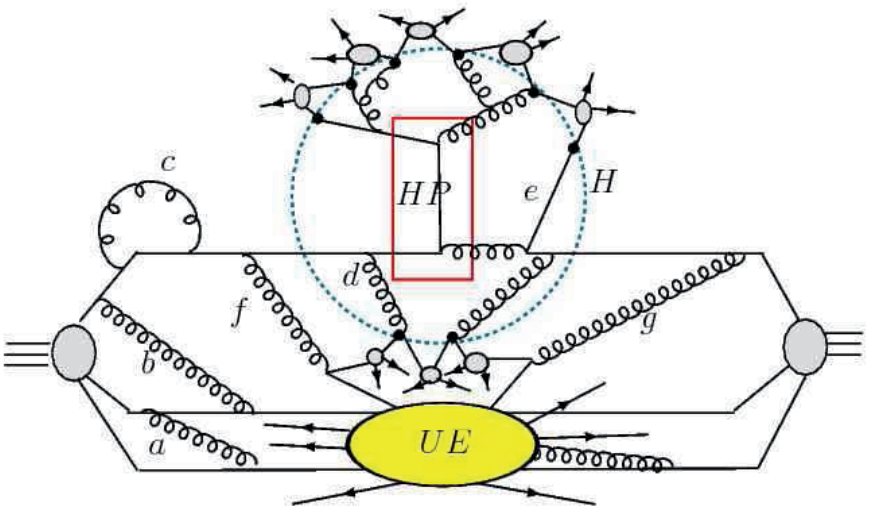

Figure 1: General structure of a hard proton–proton collision

Figure 1 illustrates how this works. As the left proton travels freely before coming into contact with the hadron coming in from the right, its constituent quarks are held together by the constant exchange of virtual gluons (e.g. gluons a and b in the picture). These gluons are mostly soft, because any hard exchange would cause the constituent

quarks to fly apart, and a second hard exchange would be necessary to reestablish the balance of momentum and keep the proton together. Gluons of high virtuality (gluon c in the picture) prefer therefore to be reabsorbed by the same quark, within a time inversely proportional to their virtuality, as prescribed by the uncertainty principle. The state of the quark is, however, left unchanged by this process. Altogether this suggests that the global state of the proton, although defined by a complex set of gluon exchanges between quarks, is nevertheless determined by interactions which have a time scale of the order of $1/m_p$. When seen in the laboratory frame where the proton is moving with energy $\sqrt{S}/2$, this time is furthermore Lorentz dilated by a factor $\gamma = \sqrt{S}/2m_p$. If we disturb a quark with a probe of virtuality $Q \gg m_p$, the time frame for this interaction is so short $(1/Q)$ that the interactions of the quark with the rest of the proton can be neglected. The struck quark cannot negotiate with its partners a coherent response to the external perturbation: it simply does not have the time to communicate to them that it is being kicked away. On this time scale, only gluons with energy of the order of Q can be emitted, something which, to happen coherently over the whole proton, is suppressed by powers of m_p/Q (this suppression characterizes the 'elastic form factor' of the proton). In the figure, the hard process is represented by the rectangle labelled HP. In this example a head-on collision with a gluon from the opposite hadron leads to a $qg \to qg$ scattering with a momentum exchange of the order of Q. This and other possible processes can be calculated from first principles in perturbative QCD, using elementary quarks and gluons as external states.

When the constituent is suddenly deflected, the partons that it had recently radiated cannot be reabsorbed (as happened to gluon c earlier) because the constituent is no longer there waiting for the partons to come back. This is the case, for example, of the gluon d emitted by the quark, and of the quark e from the opposite hadron; the emitted gluon got engaged in the hard interaction. The number of 'liberated' partons will depend on the hard scale Q: the larger the value of Q, the more sudden the deflection of the struck parton, and the fewer the partons that can reconnect before its departure (typically only partons with virtuality larger than Q).

After the hard process, the partons liberated during the evolution prior to the collision and the partons created by the hard collision will themselves emit radiation. The radiation process, governed by perturbative QCD, continues until a low virtuality scale is reached (the boundary region labelled with a dotted line, H, in our figure). To describe this perturbative evolution phase, proper care has to be taken to incorporate quantum coherence effects, which in principle connect the probabilities of radiation off different partons in the event. Once the low virtuality scale is reached, the memory of the hard-process phase has been lost, once again as a result of different time scales in the problem, and the final phase of hadronization takes over. Because of the decoupling from the hard-process phase, the hadronization is assumed to be independent of the initial hard process, and its parametrization, tuned to the observables of

some reference process, can then be used in other hard interactions (universality of hadronization). Nearby partons merge into colour-singlet clusters (the grey blobs in fig. 1), which are then decayed phenomenologically into physical hadrons. To complete the picture, we need to understand the evolution of the fragments of the initial hadrons. As shown in the figure, this evolution cannot be entirely independent of what happens in the hard event, because at least colour quantum numbers must be exchanged to guarantee the overall neutrality and conservation of baryon number. In our example, the gluons f and g, emitted early on in the perturbative evolution of the initial state, split into $q\overline{q}$ pairs which are shared between the hadron fragments (whose overall interaction is represented by the oval labelled UE, for Underlying Event) and the clusters resulting from the evolution of the initial state.

The above ideas are embodied in the following factorization formula, which represents the starting point of any theoretical analysis of cross sections and observables in hadronic collisions:

$$\frac{d\,\sigma}{d\,X} \;=\; \sum_{j,k} \int_{\hat{X}} \; f_j(x_1,Q) f_k(x_2,Q) \; \frac{d\,\hat{\sigma}_{jk}(Q)}{d\,\hat{X}} \; F(\hat{X} \to X;\, Q) \;, \tag{1}$$

where:

- X is some hadronic observable (e.g. the transverse momentum of a pion, the invariant mass of a combination of particles, etc.);
- the sum over j and k extends over the partons types inside the colliding hadrons;
- the function $f_j(x,Q)$ (known as parton distribution function, PDF) parametrizes the number density of parton type j with momentum fraction x in a proton probed at a scale Q (more later on the meaning of this scale);
- $\hat{X}$ is a parton-level kinematic variable (e.g. the transverse momentum of a parton from the hard scattering);
- $\hat{\sigma}_{jk}$ is the parton-level cross section, differential in the variable $\hat{X}$;
- $F(\hat{X} \to X;\, Q)$ is a transition function, weighting the probability that the partonic state defining $\hat{X}$ gives rise, after hadronization, to the hadronic observable X.

A simple justification of factorization, relying on some basic physics arguments, is given in the next section.

2.1 The parton densities and their evolution

As mentioned above, the binding forces responsible for the quark confinement are due to the exchange of rather soft gluons. If a quark were to exchange just a single hard virtual gluon with another quark, the recoil would tend to break the proton apart. It is easy to verify that the exchange of gluons with virtuality larger than Q is then proportional to some large power of m_p/Q, m_p being the proton mass. Since the gluon coupling constant gets smaller at large Q, exchange of hard gluons is significantly suppressed [1]. Consider in fact the picture in Fig. 2. The exchange

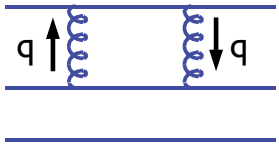

Figure 2: Gluon exchange inside the proton

of two gluons is required to ensure that the momentum exchanged after the first gluon emission is returned to the quark, and the proton maintains its structure. The contributions of hard gluons to this process can be approximated by integrating the loop over large momenta:

$$\int_Q \frac{d^4q}{q^6} \sim \frac{1}{Q^2} \,. \tag{2}$$

At large Q this contribution is suppressed by powers of $(m_p/Q)^2$, where the proton mass m_p is included as being the only dimensionful quantity available (one could use here the fundamental scale of QCD, Λ_{QCD}, but numerically this is anyway of the order of a GeV). The interactions keeping the proton together are therefore dominated by soft exchanges, with virtuality Q of the order of m_p. Owing to Heisenberg's uncertainty principle, the typical time scale of these exchanges is of the order of $1/m_p$: this is the time during which fluctuations with virtuality of the order of m_p can survive. In the laboratory system, where the proton travels with energy E, this time is Lorentz dilated to $\tau \sim \gamma/m_p = E/m_p^2$. If we probe the proton with an off-shell photon, the interaction takes place during the limited lifetime of the virtual photon, which, once more from the uncertainty principle, is given by the inverse of its virtuality. Assuming the virtuality $Q \gg m_p$, once the photon gets 'inside' the proton and meets a quark, the struck quark has no time to negotiate a coherent response with the other quarks, because the time scale for it to 'talk' to its partners is too long compared with the duration of the interaction with the photon itself. As

[1] The fact that the coupling decreases at large Q plays a fundamental role in this argument. Were this not true, the parton picture could not be used!.

a result, the struck quark has no option but to interact with the photon as if it were a free particle. Let us look in more detail at what happens during such a process. In

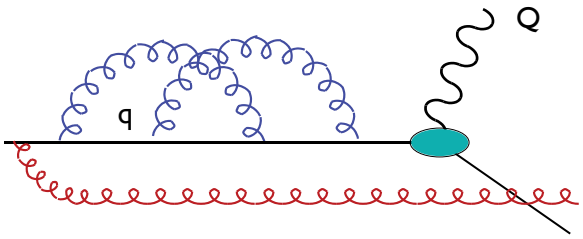

Figure 3: Gluon emission at different scales during the approach to a hard collision.

Fig. 3 we see a proton as it approaches a hard collision with a photon of virtuality Q. Gluons emitted at a scale $q > Q$ have the time to be reabsorbed, since their lifetime is very short. Their contribution to the process can be calculated in perturbative QCD, since the scale is large and in the domain where perturbative calculations are meaningful. Since after being reabsorbed the state of the quark remains the same, their only effect is an overall renormalization of the wave function, and they do not affect the quark density. A gluon emitted at a scale $q < Q$, however, has a lifetime longer than the time it takes for the quark to interact with the photon, and by the time it tries to reconnect to its parent quark, the quark has been kicked away by the photon, and is no longer there. Since the gluon has taken away some of the quark momentum, the momentum fraction x of the quark as it enters the interaction with the photon is different than the momentum it had before, and therefore its density $f(x)$ is affected. Furthermore, when the scale q is of the order of 1 GeV the state of the quark is not calculable in perturbative QCD. This state depends on the internal wave function of the proton, which perturbative QCD cannot predict. We can, however, say that the wave function of the proton, and therefore the state of the 'free' quark, are determined by the dynamics of the soft-gluon exchanges inside the proton itself. Since the time scale of this dynamics is long relative to the time scale of the photon–quark interaction, we can safely argue that the photon sees to good approximation a static snapshot of the proton's inner configuration. In other words, the state of the quark had been prepared long before the photon arrived. This also suggests that the state of the quark will not depend on the precise nature of the external probe, provided the time scale of the hard interaction is very short compared to the time it would take for the quark to readjust itself. As a result, if we could perform some measurement of the quark state using, say, a virtual-photon probe, we could then use this knowledge on the state of the quark to perform predictions for the interaction of the proton with any other probe (e.g. a virtual W or even a gluon from an opposite beam of hadrons). This is the essence of the universality of the parton distributions.

The above picture leads to an important observation. It appears in fact that the distinction between which gluons are reabsorbed and which ones are not depends on the scale Q of the hard probe. As a result, the parton density $f(x)$ appears to depend on Q. This is illustrated in Fig. 4. The gluon emitted at a scale μ has a lifetime short

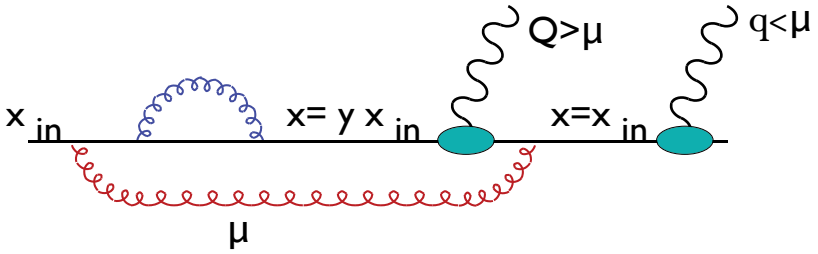

Figure 4: Scale dependence of the gluon emission during a hard collision

enough to be reabsorbed before a collision with a photon of virtuality $Q < \mu$, but too long for a photon of virtuality $Q > \mu$. When going from μ to Q, therefore, the partonic density $f(x)$ changes. We can easily describe this variation as follows:

$$f(x,Q) \;=\; f(x,\mu) + \int_x^1 dx_{in}\, f(x_{in},\mu) \int_\mu^Q dq^2 \int_0^1 dy\, \mathcal{P}(y,q^2)\, \delta(x - y x_{in})\;, \qquad (3)$$

Here we obtain the density at the scale Q by adding to $f(x)$ at the scale μ (which we label as $f(x,\mu)$) all the quarks with momentum $x_{in} > x$ that retain a proton-momentum fraction $x = y/x_{in}$ by emitting a gluon. The function $P(y,Q^2)$ describes the 'probability' that the quark emits a gluon at a scale Q, keeping a fraction y of its momentum. This function does not depend on the details of the hard process, it simply describes the radiation of a free quark subject to an interaction with virtuality Q. Since $f(x,Q)$ does not depend upon μ (μ is just used as a reference scale to construct our argument), the total derivative of the right-hand side w.r.t. μ should vanish, leading to the following equation:

$$\frac{df(x,Q)}{d\mu^2} = 0 \quad \Rightarrow \frac{df(x,\mu)}{d\mu^2} \;=\; \int_x^1 \frac{dy}{y}\, f(y,\mu)\, \mathcal{P}(x/y,\mu^2)\;. \qquad (4)$$

Dimensional analysis and the fact that the gluon emission rate is proportional to the QCD coupling squared, allow us to further write:

$$\mathcal{P}(x,Q^2) \;=\; \frac{\alpha_s}{2\pi}\,\frac{1}{Q^2}\,P(x) \qquad (5)$$

from which the Dokshitzer–Gribov–Lipatov–Altarelli–Parisi (DGLAP) equation follows:

$$\frac{df(x,\mu)}{d\log\mu^2} \;=\; \frac{\alpha_s}{2\pi}\int_x^1 \frac{dy}{y}\, f(y,\mu)\, P_{qq}(x/y)\;. \qquad (6)$$

The so-called *splitting function* $P_{qq}(x)$ can be calculated in perturbative QCD. The subscript qq is a labelling convention indicating that x refers to the momentum fraction retained by a quark after emission of a gluon.

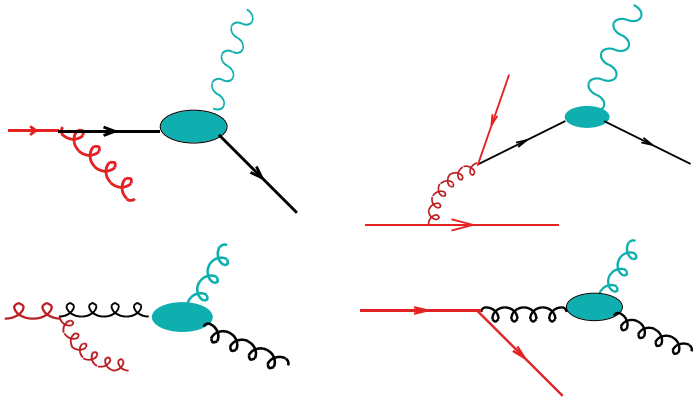

Figure 5: Top (bottom) row: the processes leading to the evolution of the quark (gluon) density

More generally, one should consider additional processes. For example, one should include cases in which the quark interacting with the photon comes from the splitting of a gluon. This is shown in Fig. 5: the upper left diagram is the one we considered above; the upper right diagram corresponds to processes where an emitted gluon has the time to split into a $q\bar{q}$ pair, and it is one of these quarks which interacts with the photon. The overall evolution equation, including the effect of gluon splitting, is given by

$$\frac{dq(x,Q)}{dt} = \frac{\alpha_s}{2\pi}\int_x^1 \frac{dy}{y}\left[q(y,Q)\,P_{qq}(\frac{x}{y}) + g(y,Q)\,P_{qg}(\frac{x}{y})\right], \tag{7}$$

where $t = \log Q^2$. For external probes that couple to gluons (for example an external gluon, coming e.g. from an incoming proton), we have a similar evolution of the gluon density (see lower row of Fig. 5):

$$\frac{dg(x,Q)}{dt} = \frac{\alpha_s}{2\pi}\int_x^1 \frac{dy}{y}\left[g(y,Q)\,P_{gg}(\frac{x}{y}) + \sum_{q,\bar{q}} q(y,Q)\,P_{gq}(\frac{x}{y})\right]. \tag{8}$$

Explicit expressions of the splitting functions $P_{ij}(x)$, as well as examples of the phenomenological consequences of the evolution equations, can be found in [4].

2.2 The final-state evolution of quarks and gluons

We discussed in the previous section the initial-state evolution of quarks and gluons as the proton approaches the hard collision. We study here how quarks and gluons evolve

after emerging from the hard process, and finally transform into hadrons, neutralizing their colours. We consider the simplest case: e^+e^- collisions, which provide the cleanest environment where to study applications of QCD at high energy. This is the place where theoretical calculations have today reached their best accuracy, and where experimental data are the most precise, especially thanks to the huge statistics accumulated by LEP, LEP2 and SLC. The key process is the annihilation of the e^+e^- pair into a virtual photon or Z^0 boson, which will subsequently decay to a $q\overline{q}$ pair. e^+e^- collisions have therefore the big advantage of providing an almost point-like source of quark pairs, so that, in contrast to the case of interactions involving hadrons in the initial state, we at least know very precisely the state of the quarks at the beginning of the interaction process.

Nevertheless, it is by no means obvious that this information is sufficient to predict the properties of the hadronic final state. We know that this final state is clearly not simply a $q\overline{q}$ pair, but some high-multiplicity set of hadrons. For example, as shown in Fig. 6, the average multiplicity of charged hadrons in the decay of a Z^0 is approximately 20. It is therefore not obvious that a calculation done using the simple

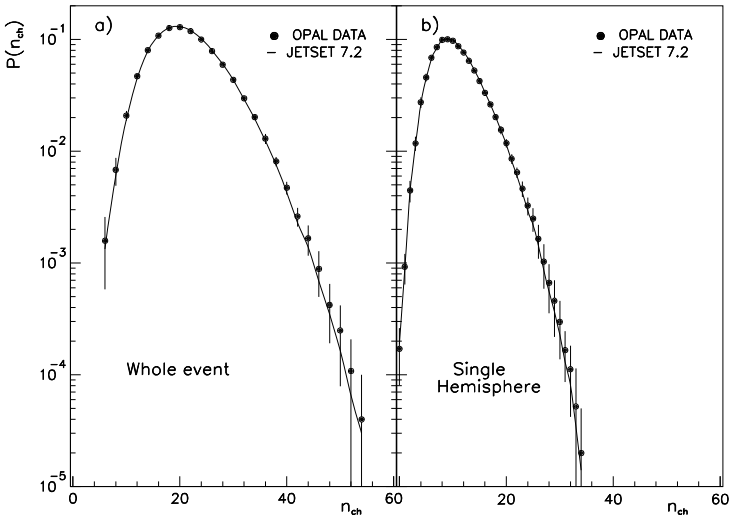

Figure 6: Charged particle multiplicity distribution in Z^0 decays

picture $e^+e^- \to q\overline{q}$ (see Fig. 7) has anything to do with reality. For example, one may wonder why we do not need to calculate $\sigma(e^+e^- \to q\overline{q}g\ldots g\ldots)$ for all possible gluon multiplicities to get an accurate estimate of $\sigma(e^+e^- \to \text{hadrons})$. And since in any case the final state is not made of q's and g's, but of π's, K's, ρ's, etc., why would $\sigma(e^+e^- \to q\overline{q}g\ldots g)$ be enough?

The solution to this puzzle lies both in a question of time and energy scales, and in the dynamics of QCD. When the $q\overline{q}$ pair is produced, the force binding q and $\overline{q}$ is

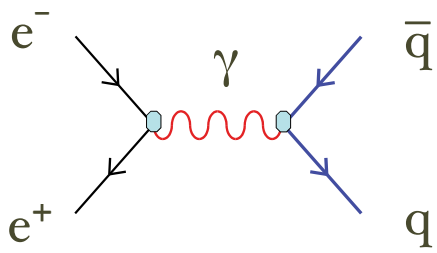

Figure 7: Tree level production of a $q\bar{q}$ pair in e^+e^- collisions

proportional to $\alpha_s(s)$ ($\sqrt{s}$ being the e^+e^- centre-of-mass energy). Therefore it is weak, and q and $\bar{q}$ behave to good approximation like free particles. The radiation emitted in the first instants after the pair creation is also perturbative, and it will stay so until a time after creation of the order of $(1\ \mathrm{GeV})^{-1}$, when radiation with wavelengths $\gtrsim (1\ \mathrm{GeV})^{-1}$ starts being emitted. At this scale the coupling constant is large, and non-perturbative phenomena and hadronization start playing a rôle. However, it can be shown that colour emission during the perturbative evolution organizes itself in such a way as to form colour-neutral, low-mass, parton clusters highly localized in phase-space. As a result, the complete colour-neutralization (i.e., the hadronization) does not involve long-range interactions between partons far away in phase-space. This is very important, because the forces acting among coloured objects at this time scale would be huge. If the perturbative evolution were to separate far apart colour-singlet $q\bar{q}$ pairs, the final-state interactions taking place during the hadronization phase would totally upset the structure of the final state.

In this picture, the identification of the perturbative cross section $\sigma(e^+e^- \to q\bar{q})$ with observable, high-multiplicity hadronic final states is realised by jets, namely collimated streams of hadrons that are the final result of the perturbative and non-perturbative evolution of each quark. The large multiplicity of the final states, shown in fig. 6, corresponds to the many particles that emerge from the collinear emissions of many gluons from each quark. The dynamics of these emissions leads these particles to grossly follow the direction of the primary quark, and the emergent bundle, the jet, inherits the kinematics of the initial quark. This is shown in the left image of fig. 8. Three-jet events, shown in the right image of the figure, arise from the $O(\alpha_s)$ corrections to the tree-level process, namely to diagrams such as those shown in fig. 9.

An important additional result of this 'pre-confining' evolution, is that the memory of where the local colour-neutral clusters came from is totally lost. So we expect the properties of hadronization to be universal: a model that describes hadronization at a given energy will work equally well at some other energy. Furthermore, so much time has passed since the original $q\bar{q}$ creation that the hadronization phase cannot

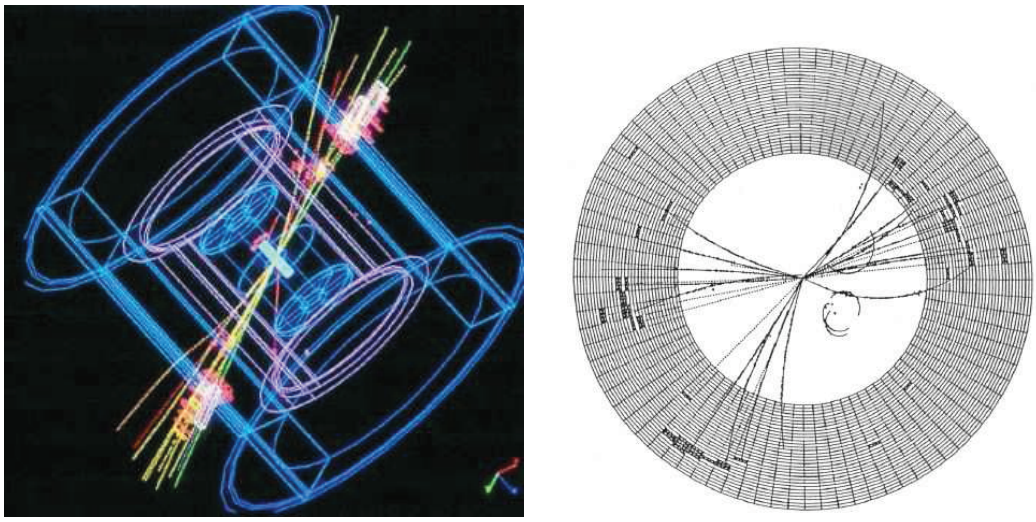

Figure 8: Experimental pictures of 2- and 3-jet final states from e^+e^- collisions

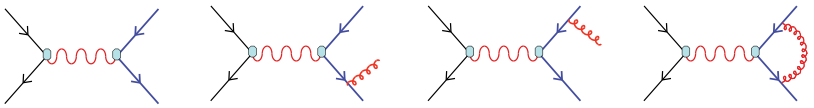

Figure 9: $O(\alpha_s)$ corrections to the tree-level $e^+e^- \to q\overline{q}$ process

significantly affect the total hadron production rate. Perturbative corrections due to the emission of the first hard partons should be calculable in PT, providing a finite, meaningful cross section.

As an example of application of this framework, we consider the calculation of the transverse energy distribution of jets produced at the Tevatron $p\overline{p}$ collider, at $\sqrt{S} = 1.96$ TeV. This is obtained by convoluting the parton densities with the partonic cross sections for gg, qg and $qq/\overline{q}q$ scattering, at next-to-leading order in QCD. The effects of final-state radiation and hadronization are estimated assuming the factorization theorem. The comparison between theory and the data of the CDF experiment [6] is shown in Fig. 10. The agreement is excellent over 8 orders of magnitude of cross section, from $E_T \sim 50$ GeV to $E_T \sim 600$ GeV. This bodes well for the reliability of theoretical calculations at the LHC!

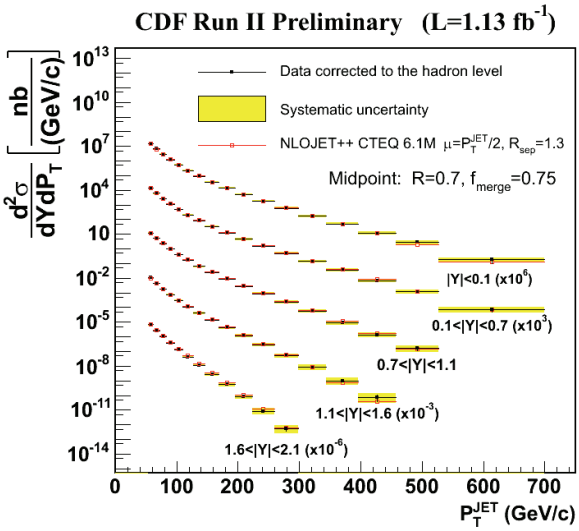

Figure 10: Inclusive jet E_T spectra, measured for various rapidity ranges by the CDF experiment at the Tevatron [6], compared to NLO QCD calculations

3 The LHC accelerator

The LHC injector chain is shown in fig. 11. The first stage of the acceleration takes place in the Linac2, a linear accelerator with an output proton energy of 50 MeV. The proton-booster synchrotron (PSB) increases the energy to 1.4 GeV, injecting into the 50-years old proton synchrotron (PS). This accelerates the beam to 26 GeV, and injects into the super proton synchrotron (SPS), out of which 450 GeV protons are eventually injected into the LHC for the start of the ramp up to the energy of 7 TeV. This is the energy acquired by a proton moving across the potential difference generated by a tower of 5 trillion household 1.5V batteries stretching half way out to the Sun! At this energy, the proton velocity is within 10^{-8} of the speed of light. Protons are grouped in 2808 bunches per beam, each bunch containing up to 10^{11} protons. While the energy of each proton is just equal to the kinetic energy of a mosquito, the full beam at nominal luminosity stores about 350 MJ of energy, like a train running at full speed.

The beam is bent along the circular LHC path by 1232 superconducting dipoles and controlled and focused by another 600 smaller magnets. The design and construction of the dipoles was the most technologically challenging part of the accelerator. To achieve the required bending power, the dipoles' field should be on average $B \sim 8.3$ Tesla. The coils are made of NiTi superconducting cable, kept at T=1.9K^o by superfluid liquid He. They are 15m long, weigh 35 tonnes and store in their magnetic field 7 MJ of energy, for a total of $\sim$10 GJ in the full ring.

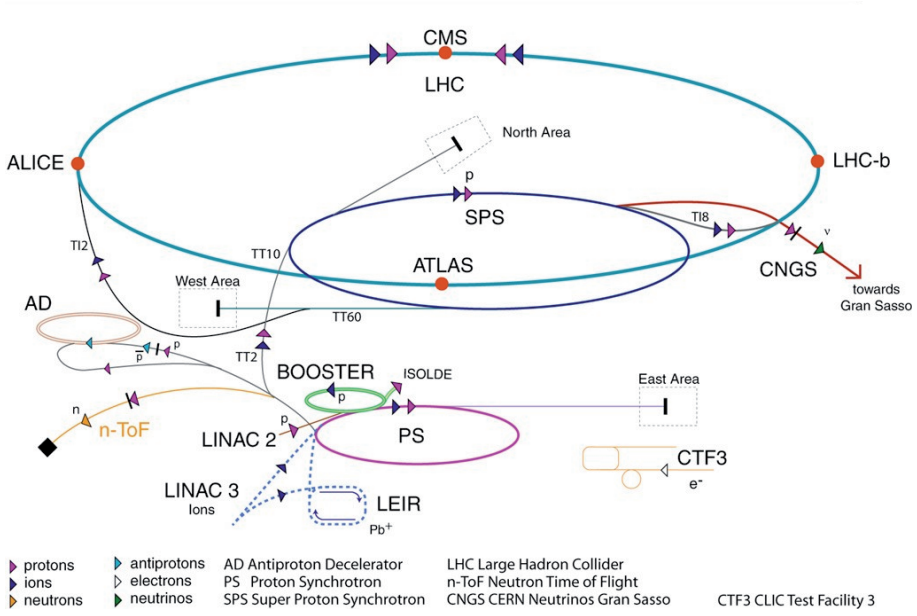

Figure 11: Layout of the full CERN accelerator complex, including all elements of the LHC injector chain. The four interaction regions hosting the main LHC experiments, ALICE, ATLAS, CMS and LHCb, are also shown.

At the collision points of the two largest experiments the bunches are squeezed to a transverse size of few tens of μm, to achieve a maximum instantaneous luminosity up to 10^{34} cm^{-2}s^{-1}, corresponding to 10 events/s for a process with cross section of 1 nb (such as the production of top quarks). The total pp inelastic cross section is of the order of 100 mb, and the number of expected inelastic pp collisions is therefore about 10^9/s. The bunch collision frequency is 40MHz (one bunch crossing each 25 ns), which implies an average of 25 pp interactions taking place simultaneously during the same bunch crossing. The amount of information that an LHC detector outputs to describe a full event (containing all of the simultaneous pp interactions) is about 1MB. The typical rate at which this information can be written to storage is a few 100MB/s, which means that only a fraction of the order of 10^{-5} of all events can be stored. Very fast and effective trigger systems are therefore required to promptly analyze online all events and select for storage the few 100 or so per second whose properties suggest they might contain interesting signals (like top quarks, Higgs bosons, supersymmetric particles, etc).

No surprise that these are the most complex experimental devices ever conceived and built in the history of science!

4 Status of the programme and timeline

The LHC has started taking data, at a reduced CM energy, at the end of November 2009. By December 2009 it has achieved pp collisions at $\sqrt{S} = 2.38$ TeV, allowing the detailed commissioning of the detectors, and the first studies of the global properties of the final states. The CM energy is planned to increase to 7 TeV by the end of February 2010, and to 10 TeV by the Summer 2010. The final jump to the design energy of 14 TeV will require hardware modifications during a year-long shut down. It is yet to be decided when this will happen, whether at the end of the currently planned run of 2010, or after further running at 10 TeV beyond 2010.

A summary of the expectations for the first two years of measurements at the LHC can be found in [8]. The timeline for the searches of the SM Higgs boson is outlined in the left plot of fig. 12, taken from [7]. This shows the amount of data,

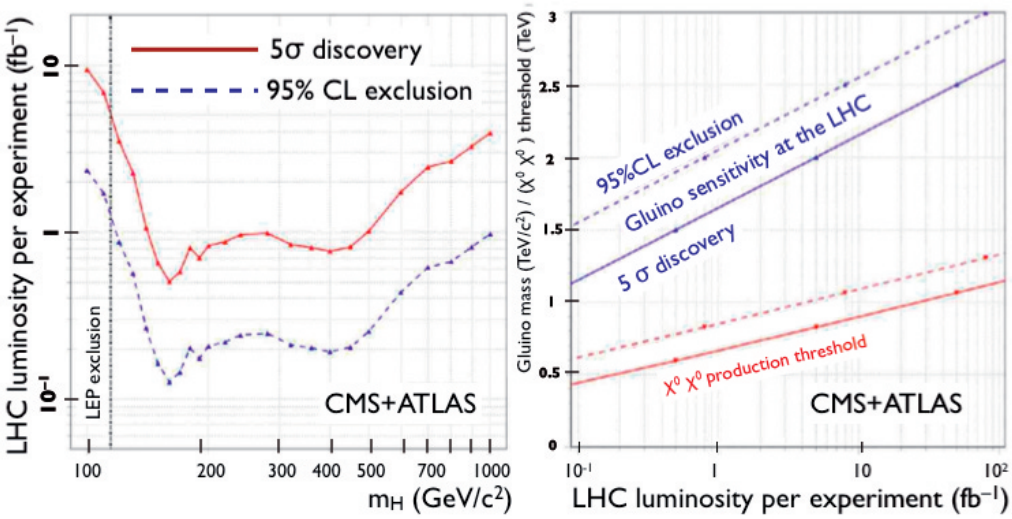

Figure 12: Discovery reach and exclusion limits at the LHC for a SM Higgs (left) and for gluinos in a supersymmetric theory (right), as a function of the integrated luminosity [7].

given in terms of *integrated luminosity*, needed by each of the two experiments to establish a 5 σ discovery, or a 95%CL exclusion, as a function of the Higgs mass[2]. The structure of dips and peaks in the plot reflects the variation of Higgs decay branching ratios as a function of the mass, defining the opening or closure of experimentally accessible final states. The present planning of LHC operations foresees the delivery of a few 100 pb^{-1} of data in 2010, and of the order of 1–few fb^{-1} over the next

[2]These numbers are given assuming collisions at 14 TeV. At 10 TeV the amount of required luminosity is larger, by a factor ~ 2 in the mass region below 200 GeV, and increasing beyond that.

couple of years at a luminosity of about 10^{33} cm^{-2}s^{-1}. After reaching the *nominal* luminosity of 10^{34} cm^{-2}s^{-1}, the LHC should start delivering about 60 fb^{-1}per year, when accounting for down-time and running efficiency.

A comparison with fig. 12 therefore shows that, within 2–3 years of data taking, the SM Higgs boson will be discovered, or entirely excluded, over the full mass range. In either case, this will signal the beginning, rather than the completion, of the LHC physics programme. Should the Higgs boson be found, an extensive campaign of studies of its detailed properties will be required, to confirm that they match the expectations of the SM, or to detect possibly minor deviations, unveiling a framework for EWSB more elaborate than the SM one. If the Higgs boson is not found, a radical departure from the SM will be needed, and the searches to understand what other mechanism is responsible for EWSB will begin.

A similar timeline is shown for the search of supersymmetry in the multijet+missing transverse energy channel, in the right plot of the figure. Once again, only 2-3 years of data taking should push the discovery reach well beyond the current limits, into the territory where one may well expect to find a positive signal, should supersymmetry exist.

Acknowledgements

I am happy to thank my friend Dima and all our hosts, for the excellent hospitality during this school, in the most beautiful of all places! The discussions with the students made the school very lively and stimulating, and I wish all of them the brightest future in physics!

Bibliography

[1] R. Barbieri, *Searching for new physics at future accelerators,* arXiv:hep-ph/0410223.

[2] G. Aad *et al.* [ATLAS Collaboration], *The ATLAS Experiment at the CERN Large Hadron Collider,* JINST **3** (2008) S08003.

[3] The CMS collaboration, *Physics Technical Design Reports,* at http://cmsdoc.cern.ch/cms/cpt/tdr/index.html.

[4] J. M. Campbell, J. W. Huston and W. J. Stirling, Rep. Prog. Phys. **70** (2007) 89 [arXiv:hep-ph/0611148].

[5] M.E. Peskin and D.V. Schroeder, An Introduction to Quantum Field Theory (Addison-Wesley, Reading, MA, 1995).

[6] T. Aaltonen *et al.* [CDF Collaboration], Phys. Rev. D **78** (2008) 052006 [arXiv:0807.2204 [hep-ex]].

[7] J. J. Blaising et al., *Potential LHC contributions to Europe's future strategy at the high-energy frontier,*, available at http://cern.ch/council-strategygroup/BB2/contributions/Blaising2.pdf.

[8] F. Gianotti and M. L. Mangano, *LHC physics: The first one-two year(s),* arXiv:hep-ph/0504221.