

1

Introduction

The theory of the strong interaction of hadrons – quantum chromodynamics, or QCD – is in many ways the most perfect and non-trivial of the established microscopic theories of physics. It is, as far as is known, a self-consistent relativistic quantum field theory. But, unlike the case of the electromagnetic and weak interactions, many primary phenomena governed by QCD are not amenable to direct calculation by weak-coupling perturbation theory. Moreover, QCD has few parameters.

To understand these assertions, first recall the classification of known microscopic interactions into strong, electromagnetic, weak, and gravitational. Precisely because the strong interaction is strong, it is useful to study QCD by itself, the other interactions being perturbations.

QCD is a quantum field theory of the kind called a non-abelian gauge theory (or a Yang-Mills theory). It has two types of field: quark fields and the gluon field. Particles corresponding to the quark fields form the basic constituents of hadrons, like the proton, with the gluon field providing the binding between quarks. There appear to be no states for isolated quarks and gluons; these particles are always confined into hadrons. This contrasts with quantum electrodynamics (QED), where instead of quarks and gluons, we have electrons and photons, which do exist in isolated single-particle states.

One key feature of QCD is “asymptotic freedom”: the effective coupling of QCD goes to zero at zero distance. Thus short-distance processes yield to the highly developed methods of Feynman perturbation theory. Among other things, this allows a perturbative analysis to give a correct renormalization of the ultra-violet divergences of QCD. The theory therefore exists in a way that the electroweak theory may not. Thus QCD contains no hints of its own breakdown.

On the other hand, unlike the case of QED, where perturbative methods give (first principles) predictions of spectacular accuracy, many apparently simple phenomena in QCD are difficult and non-perturbative, for example, its simplest bound states, like the pion and proton. Although Monte-Carlo lattice calculations have made enormous progress, they are limited both in achievable accuracy and in the observables that can be predicted, and these include no hadronic high-energy scattering processes, for example. So QCD is highly non-trivial.

Moreover, its consequences are enormous. Of course, QCD underlies the whole of nuclear physics and it creates most of the mass of ordinary matter, as contained in the proton and neutron.

Despite the non-perturbative nature of the particle states in QCD, there is a vast domain where perturbative methods can actually be applied to realistic scattering processes in QCD. The purpose of this book is to give a systematic account of these methods and their justification.

At present these are the methods that show the power of QCD most strongly. They have an almost universal impact on experiments in high-energy physics, particular with hadron beams. This ranges from the long-range planning of experiments to the analysis of data, even when the primary subject of study is a non-QCD phenomenon, such as the weak interaction, the Higgs boson, supersymmetry, etc.

It is easy to state the characteristic method, hard-scattering factorization, that enables perturbative QCD to be systematically applied to these reactions. As an example, consider production of the predicted Higgs boson in proton-proton collisions at the Large Hadron Collider (LHC), which at the time of completion of this book (2010) was starting operation at the European Organization for Nuclear Research (CERN). In the factorization approach, the proton beams are treated as collections of so-called partons: quasi-free quarks and gluons, whose (non-perturbative) distributions, the parton densities, have been measured in other experiments, and are used at the LHC with the aid of the perturbative evolution equations of Dokshitzer, Gribov, Lipatov, Altarelli, and Parisi (DGLAP). The cross sections for collisions of partons of the various types to make the Higgs boson are calculated perturbatively given its expected couplings. A provable consequence of QCD is that many useful physical cross sections are predicted by convoluting measured parton densities with short-distance partonic cross sections. It is the proved universality (or, more generally, modified universality) of the non-perturbative parton densities, etc., between different experiments that gives perturbative QCD its predictive power.

Among other things, factorization allows an extrapolation of the physics by an order of magnitude or more in center-of-mass energy from previous experiments. Many cross sections of interest are so minute (down to femtobarns at the LHC) compared with the total cross section (close to 100 mb) that considerable quantitative understanding of QCD physics is necessary for a good analysis of experimental data.

There are many places where one can learn *how* to perform perturbative QCD calculations. But a newcomer or an outsider can be forgiven for questioning the logical foundations of the subject. For example, why should calculations be made with on-shell massless partons, as is commonly done? There seems to be an essential use of actual collinear singularities associated with on-shell massless partons; how can these be so routinely and cavalierly manipulated when we know that quarks and gluons are confined inside hadrons and are clearly not free particles and are therefore definitely not on-shell?

Therefore the purpose of this book is to give a connected logical account of the methods of perturbative QCD. The intended audience includes not only graduate students in high-energy physics, but also established researchers, both in high-energy physics and elsewhere, who want a clear account of the subject.

Readers are assumed to have a knowledge¹ of relativistic quantum field theory, up to non-abelian gauge theories and the elements of renormalization theory, together with a basic knowledge of elementary particle physics. Beyond this I try to keep the treatment self-contained.

There is a clear danger that the treatment gets bogged down in mathematical minutiae without getting to the practically applicable meat of the subject. But without a sufficiently clear and precise treatment, the concepts get muddled, further development is stymied, and the construction of new innovative and correct concepts is hindered.

Indeed, not everything in perturbative QCD is properly clear and established. One reason for such problems is the way in which much knowledge in perturbative QCD has been constructed.² It is common in science to induce theoretical ideas from a pattern perceived in a body of experimental data. But in QCD, we also often induce new higher-level theoretical ideas from calculations within QCD. For example, on the basis of a set of Feynman-graph calculations, one might see a pattern that can be formalized in the statement of a factorization property together with a definition of parton densities. The induced property can be tested both by further theoretical calculations and by comparison of its predictions with experiment.

Such a factorization property has the form of a statement of a mathematical theorem, and the soundest method of establishing the property is by proving it mathematically. Naturally researchers try to do so. But because of the difficulty of perturbative QCD, there are often interesting gaps in the proofs.

The theorems of perturbative QCD are supported not only by proofs, but also by a combination of agreement with the results of particular Feynman-graph calculations and agreement with experimental data. So a gap in a proof does not imply that a theorem is actually wrong. But the gaps can be frustrating to a newcomer learning the subject. They are suggestive of things that are difficult and not fully explicitly understood; the understanding in the collective consciousness of the workers in the field is quite non-verbal.³ Such gaps could become particularly important in generalizations of the theorems.

In this book I try to make the gaps explicit. I will point out some of the danger areas, and suggest targets for research. I was also able to fill in or reduce many of the gaps.

1.1 Factorization and high-energy collisions

Since the idea of hard-scattering factorization is so central to the applications of QCD, it is useful at this point to formulate it quantitatively in a particular example. The results in this section are stated without any attempt at justification, the aim being to give a hint of the landscape we will explore in detail in the rest of the book. The section may be somewhat mysterious to a reader without any exposure to the general subject matter, and it can be skipped if necessary.

¹ Standard references include: Sterman (1993); Peskin and Schroeder (1995); Weinberg (1995, 1996); Srednicki (2007).

² These issues actually apply more generally in quantum field theory and in high-energy theory.

³ A classic case of a similar situation is with Dirac's delta function.




Fig. 1.1. (a) Drell-Yan cross section at lowest order in hard scattering. This represents an amplitude times its complex conjugate, with a sum and integral over the final state, at the vertical line. (b) Same as diagram (a) with interactions between the hadron remnants that fill in the gap between the remnants.

Let us choose the Drell-Yan process, the inclusive production of a high-mass muon pair, $H_A H_B \rightarrow \mu^+ \mu^- X$. Here H_A and H_B are incoming hadrons of momenta P_A and P_B in a collision with a large center-of-mass energy \sqrt{s} . The muon pair is produced through a virtual photon (or other electroweak boson).⁴ The symbol X denotes that the rest of the final state is summed/integrated over, and is treated as unobserved.

In factorization, the $\mu^+ \mu^-$ pair is formed in the interaction of one constituent out of each hadron, with the lowest-order case being simple quark-antiquark annihilation, as in Fig. 1.1(a). In the center of the figure, a quark out of one hadron and an antiquark out of the other annihilate at a vertex with a line for a highly virtual photon. At the other end of the photon line are the detected muon and antimuon. The shaded blobs represent the remnants of the incoming hadrons.

Factorization applies when the $\mu^+ \mu^-$ pair has high mass. In this context the constituents are called partons, and the possible types correspond to the fields in QCD: the quarks of various types (called “flavors”), and gluons.

In a sense to be made precise in Ch. 6, the partons are approximately aligned with their parent hadrons, and we will define a momentum fraction ξ of a parton with respect to its parent. We will define “parton densities”, $f_{i/H}(\xi)$. Here i labels the type of parton, and H the parent hadron. Then $f_{i/H}(\xi)$ is treated as the number density of partons of type i in hadron H . The factorized cross section is

$$\frac{d\sigma}{dQ^2 dy} = \sum_{ij} \int_0^1 d\xi_a \int_0^1 d\xi_b f_{i/H_A}(\xi_a) f_{j/H_B}(\xi_b) \frac{d\hat{\sigma}(\xi_a, \xi_b, i, j)}{dQ^2 dy}, \quad (1.1)$$

where y is the rapidity⁵ of the $\mu^+ \mu^-$ pair, and Q is its invariant mass. We have chosen to integrate over q_T , the transverse momentum of the pair relative to the collision axis. The errors in the factorization formula are suppressed by a power of a hadronic mass divided by Q or \sqrt{s} . We have a sum over the types of the parton that are involved, one out of

⁴ Any other type of lepton pair may be similarly treated, e.g., $e^+ e^-$ or $\mu^+ \nu_\mu$.

⁵ See Sec. B.3 for a definition.

each beam hadron, and we have an integral over the possible momentum fractions. The partons themselves undergo a collision, and this gives the $\mu^+\mu^-$ pair, with an effective cross section denoted by $d\hat{\sigma}(\xi_a, \xi_b, i, j)/dQ^2 dy$. This we call a “short-distance cross section” or a “hard-scattering coefficient”. It differs conceptually from an ordinary cross section in that it is arranged to be not a complete cross section for partonic scattering, but to contain only short-distance contributions.

The short-distance partonic cross section can be usefully calculated in powers of the coupling $\alpha_s(Q)$, which is small when Q is large, because of the asymptotic freedom of QCD. The lowest-order hard-scattering for the Drell-Yan process is the tree diagram for quark-antiquark annihilation to $\mu^+\mu^-$, at the center of Fig. 1.1(a). At this order, the parton momentum fractions are determined from the muon-pair kinematics: $\xi_a = e^y Q/\sqrt{s}$, $\xi_b = e^{-y} Q/\sqrt{s}$, in the center-of-mass frame. Thus the Drell-Yan cross section gives a direct probe of the underlying quark and antiquark.

The correct definition of a parton density depends on a resolution scale, which can be usefully set equal to Q . There is an equation, the DGLAP (Dokshitzer-Gribov-Lipatov-Altarelli-Parisi) equation, which governs the scale dependence. It is a linear integro-differential equation – (8.30) below – whose kernel is also perturbatively computable.

There are many similar factorization formulae for a wide variety of processes, typified by the production of a high-mass electroweak system ($\mu^+\mu^-$ pair, and W , Z and Higgs bosons), or by the production of particles with large transverse momentum. There are, in addition, many further developments, characterized by more complicated kinematic situations.

We can now appreciate some of the issues that make the derivation and understanding of factorization very non-trivial. One is that in Fig. 1.1(a) there are beam remnants going into the final state. We can treat each of these as each being boosted from the rest frame of its parent hadron to a high energy, so that the final state appears to have two oppositely highly boosted systems, with a distinctive large rapidity gap between them. Such a state is in fact sometimes seen experimentally, and is called a “diffractive” configuration. But diffractive events are only a small fraction of the total. Therefore experiment suggests that the beam remnants interact with high probability, as notated in Fig. 1.1(b). We will need to understand whether or not factorization gets violated. For the fully inclusive Drell-Yan process, we will show that a quite non-trivial cancellation applies, so that factorization actually works: Ch. 14.

Moreover, the beam remnants in Fig. 1.1(a) appear to have quantum numbers, e.g., fractional electric charge, that correspond to beam hadrons with a quark or antiquark removed. This is not observed in actual diffractive events.

Another issue that we will solve is that the colliding partons are often treated as being free and on-shell, even though they are bound inside their hadrons. In addition, the partons are treated as having certain fractions of their parent hadrons’ momenta, even though the parton and hadron momenta cannot always be exactly parallel; there is certainly a distribution over the components of parton momentum transverse to the collision axis. We will see how this is allowed for in defining parton densities, and how in some situations we need to treat it more explicitly. One standard example is the Drell-Yan process when we take the cross

section differential in the transverse momentum q_T of the $\mu^+\mu^-$ pair instead of integrating over it as we did in (1.1).

1.2 Why we trust QCD is correct

QCD is generally regarded as the correct theory of the strong interaction. The reasons are not just that it makes successful quantitative predictions. Also important are structural arguments, as summarized in Ch. 2. These arguments start from high-level abstractions from data, and make a quite rigid deduction of QCD as the theory of the strong interaction. There does not appear to be an underdetermination of theory by data. The arguments are like those that led Einstein to the theories of special and general relativity, and to those that led Heisenberg, Born, Jordan, and Dirac to founding the current formulation of quantum theory. Once QCD and asymptotic freedom were discovered, there was a positive feedback loop where successful experimental predictions confirmed QCD. But the structural arguments are also critical in giving the realization that QCD applies generally to strong-interaction phenomena, including those that not yet derived from QCD, e.g., the bulk of conventional nuclear physics.

As with many modern theories of physics, QCD and its applications have many elements that do not in any immediate sense correspond to tangible real-world entities. Parton densities are a good example. These elements, and the associated mathematics, provide links between many different kinds of experimental data.

Indeed it is in the links between experiments that many of the predictions of QCD arise. Perturbative calculations alone do not predict any cross sections in hadron-hadron scattering and lepton-hadron scattering. Factorization gives us cross sections in terms of non-perturbative parton densities, which we currently cannot compute from first principles. But from QCD we deduce that the parton densities are universal between different reactions, with energy-dependent modifications of universality caused by the evolution of parton densities with a scale parameter. So we can fit parton densities from some limited set of data, and use them to predict other cross sections, with the aid of perturbatively calculated coefficients like $d\hat{\sigma}$ in (1.1).

No single experiment really provides a critical test of QCD. But a sufficiently large collection of experiments does simultaneously provide measurements of non-perturbative factors, tests of the factorization structure, and tests of QCD itself. There are some interesting issues in statistics and in the philosophy of science here, which do not appear to arise in such a strong form in other areas of science.

1.3 Notation

As regards normalization conventions and the like, I generally follow the conventions of the Particle Data Group (PDG) (Amsler *et al.*, 2008), since this forms a standard for our field. I point out exceptions explicitly. For a collection of many notations and standard results, see App. A. For acronyms and abbreviations, see Sec. A.3.

For the most part, I use the “natural units” conventional to the field, where $\hbar = c = \epsilon_0 = 1$, and the GeV is used as the unit of energy. See Sec. A.2 for common conversion factors.

I have found some variations on the standard symbol for equality to be useful. First, to flag the definition of something, I put “def” over the “=”, as in

$$Q \stackrel{\text{def}}{=} \sqrt{-q^2}. \quad (1.2)$$

Quite often, it is convenient to explain the definition of some conceptually difficult object by first proposing a simplified candidate definition based on a naive picture of the physics, and then building up to the correct definition, e.g., with parton densities, starting in Ch. 6. So I use the notation

$$\text{Quantity} \stackrel{\text{prelim}}{=} \text{preliminary candidate definition}, \quad (1.3)$$

for the early incorrect definitions. This avoids having several apparently incompatible definitions of the same quantity, with the wrong ones prone to being taken out of context.

Similarly there are situations where I motivate proofs and statements of difficult results by first formulating them in a simplified situation. For example, before formulating and deriving true factorization theorems, I examine the parton model, an intuitive approximation to real physics that is relatively easy to motivate. So I flag these suggestive but ultimately incorrect results with a question mark over an equality sign, e.g.,

$$F_L \stackrel{?}{=} 0. \quad (1.4)$$

1.4 Problems and exercises

I have devised a number of chapter-end exercises. Some of these are relatively elementary and should be tackled by anyone learning the subject. Among these are exercises to complete derivations in the text; some others explore the conceptual framework by the derivation of further results. There are also harder exercises, rated with one to five stars. Those with three or more stars are really research problems. I do not necessarily have any answers or even suggestions for approaches to the research problems; a good solution could easily be suitable for journal publication.

2

Why QCD?

A possible approach to a theory like QCD is just to state its definition, and then immediately proceed deductively. However, this begs the question of why we should use this theory and not some other. Moreover, the approach is quite abstract, and the initial connection to the real physical world is missing.

Instead, I will take a quasi-historical approach, after first stating the theory. Such an approach is suitable for newcomers since their background in QCD is like that of its inventors/discoverers, i.e., little or none. There were several lines of development, all of which powerfully converged on a unique theory from key aspects of experimental data. Of course, we see this much more readily in retrospect than was apparent at the time of the original work, and my account is selective in focusing on the issues now seen to be the most significant. A historical approach also enables the introduction of ideas and methods that do not specifically depend on QCD: e.g., deeply inelastic scattering and the parton model.

I have tried to make the presentation self-contained, in summarizing the relevant experimental phenomena and their consequences for theory. The reader is only assumed to have a working knowledge of relativistic quantum field theory. Inevitably there are issues, ideas, experiments, and historical developments which will be unfamiliar to many readers, and for which a complete treatment needs much more space. I give references for many of these. In addition, there are several references that are global to the whole chapter and that the reader should refer to for more detail. A detailed historical account from the point of view of a physicist is given in the excellent book by Pais (1986). A good account of the phenomena is given by Perkins (2000). Standard books on quantum field theory also refer to them; see, for example, Sterman (1993); Peskin and Schroeder (1995); Weinberg (1995, 1996); Srednicki (2007). A comprehensive account of experimental results is given by the Particle Data Group in Amsler *et al.* (2008); this includes up-to-date authoritative summaries of measurements and their theoretical interpretation.

Naturally, QCD is not the whole story; there are known electromagnetic, weak and gravitational interactions, and presumably if we examine phenomena at short enough distances, beyond the reach of current experimental probes, we are likely to need new theories. But within the domain of the strong interaction at accessible scales, there is an amazing uniqueness to the structure of QCD.

2.1 QCD: statement of the theory

An expert in quantum field theory could simply define QCD as a standard Yang-Mills theory with a gauge group SU(3) and several multiplets of Dirac fields in the fundamental (triplet) representation of SU(3).

In more detail, QCD is specified by its set of field variables and its Lagrangian density \mathcal{L} . The Dirac fields $\psi_{\rho af}$ are called quark fields, and the gauge fields A_μ^α are called gluon fields. On the quark fields the indices ρ , a , and f are respectively a Dirac index, a “color” index taking on three values, and a “flavor” index. The gauge group acts on the color index. Currently the flavor index has six known values u, d, s, c, b, t (or “up”, “down”, “strange”, “charm”, “bottom”, and “top”). On the gluon field, the color index α has eight values, for the generators of SU(3), and μ is a Lorentz vector index. The important role played by the color charge leads to the theory’s name, “quantum chromodynamics” or QCD. Of course, the names “color” and “flavor”, and the names of the quark flavors, are whimsical inventions unrelated to their everyday meanings.

To deal with the renormalization of the UV divergences of QCD (Sec. 3.2) we distinguish between bare and renormalized quantities (fields, coupling and masses). We define QCD by a Lagrangian written in terms of bare quantities, which are distinguished by a subscript 0 or (0). The gauge-invariant Lagrangian is the standard Yang-Mills one:

$$\mathcal{L}_{\text{GI}} = \bar{\psi}_0(i\not{D} - m_0)\psi_0 - \frac{1}{4}(G_{(0)\mu\nu}^\alpha)^2. \tag{2.1}$$

The full Lagrangian used for perturbation theory will add to this some terms to implement gauge fixing by the Faddeev-Popov method; see Sec. 3.1. The covariant derivative is given by

$$D_\mu \psi_0 \stackrel{\text{def}}{=} (\partial_\mu + ig_0 t^\alpha A_{(0)\mu}^\alpha) \psi_0, \tag{2.2}$$

where t^α are the standard generating matrices¹ of the SU(3) group, acting on the color indices of ψ . The gluon field strength tensor is

$$G_{(0)\mu\nu}^\alpha \stackrel{\text{def}}{=} \partial_\mu A_{(0)\nu}^\alpha - \partial_\nu A_{(0)\mu}^\alpha - g_0 f_{\alpha\beta\gamma} A_{(0)\mu}^\beta A_{(0)\nu}^\gamma, \tag{2.3}$$

where $f_{\alpha\beta\gamma}$ are the (fully antisymmetric) structure constants of the gauge group, defined so that $[t_\alpha, t_\beta] = if_{\alpha\beta\gamma} t_\gamma$. The Lagrangian is invariant under local (i.e., space-time-dependent) SU(3) transformations:

$$\psi_{(0)\rho af}(x) \mapsto [e^{-ig_0\omega_a(x)t^\alpha}]_{ab} \psi_{(0)\rho bf}(x), \tag{2.4a}$$

$$A_{(0)\mu}^\alpha(x)t^\alpha \mapsto \frac{-i}{g_0} e^{-ig_0\omega_a(x)t^\alpha} D_\mu e^{ig_0\omega_a(x)t^\alpha}. \tag{2.4b}$$

The quark fields have been redefined, as is always possible (Weinberg, 1973a), so that the mass matrix is diagonal:

$$\bar{\psi}_0 m_0 \psi_0 = m_{0u} \bar{u}_0 u_0 + m_{0d} \bar{d}_0 d_0 + m_{0s} \bar{s}_0 s_0 + \dots \tag{2.5}$$

¹ $t^\alpha = \frac{1}{2}\lambda^\alpha$, where the standard λ^α are given in, e.g., Amsler *et al.* (2008, p. 338).

Here separate symbols are used for the fields for different flavors of quark: $u_{0\rho a} = \psi_{0\rho au}$, etc.

The renormalized masses of the quarks are given in Table 2.2 below, along with the masses of the other elementary particles of the Standard Model. Large fractional uncertainties for the light quark masses arise because quarks are in fact confined inside color-singlet hadrons, which gives considerable complications in relating the mass parameters to data.

For their electromagnetic interactions, we need the quark charges:

$$e_d = e_s = e_b = -1/3, \quad e_u = e_c = e_t = 2/3, \quad (2.6)$$

in units of the positron charge.

The only significant freedom in specifying QCD is in the set of matter fields, the quarks. At the time of discovery of QCD, only the u , d and s quarks were known; the c quark came slightly later. The discovery of the b and t quarks needed high enough collision energies to produce them. There have been many conjectures about possible new heavy quarks, both scalar and fermion, possibly in non-triplet color representations, but searches so far have been unsuccessful (Amsler *et al.*, 2008). The decoupling theorem (Appelquist and Carazzone, 1975) for heavy fields ensures that we can ignore the heavy fields if experimental energies are too low to make the corresponding particles.

A complete theory of strong, electromagnetic, and weak interactions is made by combining QCD with the Weinberg-Salam theory to form the Standard Model of elementary particle physics, summarized in Sec. 2.7.

2.2 Development of QCD

Why we should postulate the QCD Lagrangian and study QCD as the unique field theory for the strong interaction? An answer to this question should be at a high level and broad, since QCD is a high-level theory, intended to cover a broad range of phenomena, i.e., all of the strong hadronic interaction.

Starting in the 1950s, as accelerator energies increased, elementary particle physics gradually became a separate subject, distinct from nuclear physics. Several, not entirely distinct, strands of research led to the discovery of QCD in 1972–1973:

1. The quark model of hadron states.
2. The (successful) search for a theory of the weak interactions of leptons, including the weak interactions of hadrons.
3. Current algebra, i.e., the analysis of the currents for the (approximate) flavor symmetries of the strong interaction, including their relationships to the electroweak interactions of hadrons.
4. The theoretical development of non-abelian gauge theories.
5. Deeply inelastic lepton scattering and the measurement that the strong interaction is quite weak at short distances.