Introduction to Particle Physics

Brian Shuve

Contents

These notes are intended to introduce students to many of the basic ideas in particle physics needed to begin research. They were written to be accessible to first-year physics students at Harvey Mudd College, who have taken courses in calculus and calculus-based physics, including special relativity. Even if you don't have this background, however, you can skip over the more technical bits and still get some useful insight into underlying ideas in particle physics.

1 What Is Particle Physics?

Particle physicists work to understand how the universe works on its most fundamental level. To do this, we first need to identify the basic building blocks of our universe. We then want to understand how they can be assembled in various ways to give the diverse array of objects that we see in our world: atoms, molecules, cells, people, planets, stars and galaxies. What forces act on these building blocks in order to explain the structures we see?

Consider for a moment a non-relativistic point-like object obeying Newton's Second Law:

$$\vec{F}(\vec{r}) \quad = \quad M\vec{a}, \tag{1}$$

which is the usual statement that an applied net force leads the acceleration of an object. If we specify an applied force and a numerical value for mass, then this law tells us the object's acceleration, which in turn tells us how the object's position and velocity change in the following moment. We can repeat this over and over to determine the object's complete motion over time. For an arbitrarily complicated system of forces and masses, we can apply the Second Law separately to each object, and given an initial configuration of positions and velocities, we can use it to find out where everything ends up after some period of time has elapsed[1]. To summarize, the ingredients we need to fully determine the motions of a system (whether it be a point mass, a set of atoms in a material, or everything in the entirely Universe) are:

- The masses of the constituents of the system (*i.e.*, the types of objects that we have);
- The forces that act on the constituents of the system (*e.g.,* whether an object experiences the electric force and what its electric charge is).

Once we have these two sets of properties, we can use Newton's Second Law (or the quantum equivalent) to determine how the objects move around under the various forces. Particle physics is essentially the branch of physics that tries to figure out these two things: what, deep down, are the elementary types of matter that exist in nature, and what forces give rise to interactions between them? The collection of elementary particles and forces are known collectively as the **Standard Model of particle physics.**

Let's start with the forces. In your introductory physics class, you no doubt became acquainted with a whole variety of forces: push and pull forces, gravitational forces, electric and magnetic forces, drag and friction forces, etc. Each of these has a different (and usually complicated) dependence on the objects in a system. For example, the friction force between two objects depends on whether the objects are in relative motion, the magnitude of the normal force between them, the type of material, and so on. Friction is by definition something that acts on *macroscopic* objects that are made up of many atoms and molecules. If, however, we zoom in on the interface between two objects and look at what's going on at the microscopic level, what we see deep down is that friction actually results from

[1]This is a purely classical equation, but the basic argument is true in both classical and quantum mechanics. In quantum mechanics, the force and the properties of the particle (mass, spin) are all encoded in the Hamiltonian, while the motion of the particle is determined by the evolution of the wavefunction or state vector.

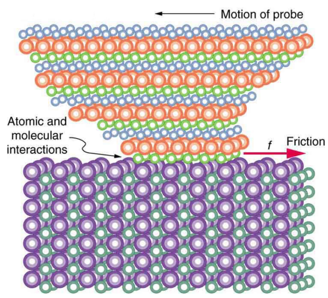

Figure 1: Illustration of how the macroscopic force of friction originates from the microscopic electrical attractions and repulsions between electrons and protons in each of the materials (image source: Rice Open CNX).

the formation and breaking of chemical bonds between atoms in each material (see Fig. 1). This "friction" force, then, is actually a phenomenon that emerges from a more fundamental interaction, namely the electrostatic attractions and repulsions between components of atoms. In fact, most of the forces I listed (the drag forces, the push and full forces) have the same property: they derive from the action of inter-atomic electric forces that exist between atoms within large objects.

As particle physicists, we are interested only in the **fundamental forces** of nature. These are the forces that are irreducible, meaning that they cannot be explained in terms of the action of more basic forces. Friction is not a fundamental force because it emerges from the complicated motions of large numbers of electrically charged particles. The electric force is a manifestation of a fundamental force, though, because it is not in turn a manifestation of anything simpler.

If we focus on fundamental forces, we don't have to worry about drag, pushing, or pulling forces like we often consider in a mechanics course. Instead, we only have to think about a small set of **four** fundamental forces:

1. **The Electromagnetic Force:** Responsible for electric and magnetic forces. In everyday life, we usually observe electric and magnetic phenomena as separate. However, the motion of electrically charged objects induces magnetic fields (see Fig. 2), and the dynamics of electric and magnetic fields are inseparable[2]. Electricity and magnetism are therefore combined into a single fundamental force known as the electromagnetic, or EM, force. The EM force explains most of the forces we experience on a day-to-day basis: for example, we don't fall through the floor because of the repulsive forces existing between electrically charged electrons and protons inside of adjacent atoms that prevents their arbitrary compression. Indeed, electromagnetism is responsible for each everyday force we experience with the exception of...

2. **The Gravitational Force:** This is a universally attractive force that exists between every type of matter and is proportional to an object's mass-energy (see Fig. 3). It is an exceptionally feeble force, and is only important to us because we happen to be sitting on top of an enormous chunk of rock (and orbiting around an even larger ball of burning gas). While gravity acts on all fundamental particles, it is usually so weak that its effects are unobservable when looking at individual atoms (and smaller particles).

3. **The Strong Force:** This force is responsible for binding sub-nuclear particles known as **quarks** into protons and neutrons. It is so strong that, except at extremely high energies, it is not

[2] Indeed, observers moving at different speeds relative to a system of electrically charged particles will disagree on the magnitudes of the electric and magnetic fields they observe, although all observers agree on how these different electric and magnetic fields cause objects to move.

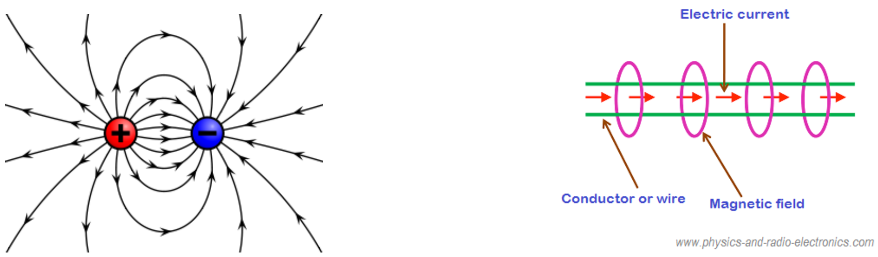

Figure 2: (Left) Electric field induced by a configuration of electrically charged particles (image source: Wikipedia). (Right) Magnetic field induced by a moving current of electric charges.

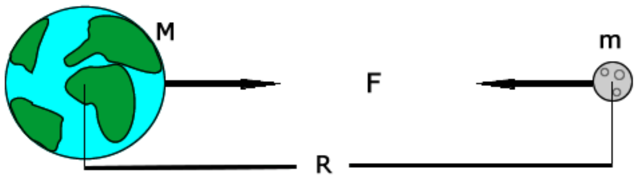

Figure 3: The gravitational force always pulls massive objects towards one another. It is also exceptionally feeble, only being important in the vicinity of very heavy objects like the Sun and Earth (image source: School for Champions).

possible to pull quarks outside of protons and neutrons; they are, for all intents and purposes, permanently bound inside of protons and neutrons. The strong force also "leaks" outside of the protons and neutrons, providing an additional attractive force between these particles. For example, this residual strong force keeps protons and neutrons stuck together to form nuclei. We know this force is *very* strong, in part because otherwise the positively charged protons inside of nuclei would repel and the electric force would blow the nucleus apart. These effects are shown in Fig. 4.

4. **The Weak Force:** In spite of the name, this fourth force is not always the weakest force (and is certainly much stronger than gravity). If we consider nuclei at rest inside of atoms, however, it is weaker than the strong and electromagnetic forces and so we are stuck with the name. Its only noticeable effect at the atomic level is that it is responsible for certain kinds of radioactive decay (such as the emission of beta particles). The beta decay process resulting from the weak force is shown in Fig. 5.

That's it. As far as we know, every phenomenon we observe in nature follows from these four forces[3]. We will explore the details of each of these forces further in Section 2.
Exercise: In the following examples, describe how the fundamental forces act to give rise to the specified motion:

- *A person standing on the floor proceeds to jump up in the air and fall back down;*
- *A kettle is placed on the stove and heated until the water inside boils.*

Now that we have the forces, we need to know the particles or types of matter on which the forces act. Hopefully, you are already familiar with the fact that everything we see in the world around us

[3]In case there is a professional high-energy theorist reading this (or you are a bit of a pedant), there is also a "Yukawa force" resulting from the interactions of the Higgs boson with fermions. In this section, though, I'm using force in the conventional particle-physics sense as a proxy for an interaction mediated by gauge bosons.

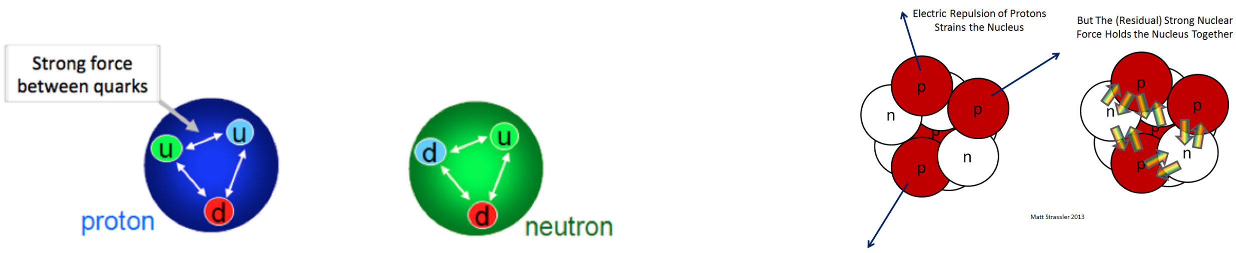

Figure 4: (Left) The strong force keeps quarks stuck together inside of protons and neutrons. (Right) The strong force can also "leak" outside of protons and neutrons, meaning that there is a residual strong force that attracts protons and neutrons together. This allows the protons and neutrons to remain stuck in an atomic nucleus, while the electrostatic force would otherwise push the protons apart (image source: Matt Strassler).

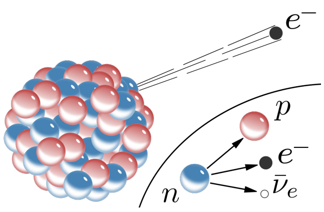

Figure 5: The weak force mediates the process known as radioactive beta decay. In this process, the weak force converts a neutron (marked as blue) into a proton (marked as red) along with an electron and an invisible particle called a neutrino. The proton remains inside the nucleus, while the electron and neutrino are energetically ejected from the nucleus (image source: Wikipedia).

(such as the desk you are sitting at or the computer you are using) is made up of atoms. These atoms, in turn, are made up of subatomic particles such as protons, neutrons, and electrons: the properties of different types of atoms are determined solely by how many electrons, protons, and neutrons are in each atom. The protons and neutrons are themselves made up of quarks; see Fig. 6. If we continue to zoom in as far as we can go, we find a collection of some twenty "fundamental" particles which, to our knowledge, are not made up of anything smaller. The particles can be summarized in a table (Fig. 7). They are:

1. **The Quarks:** These are the particles that bind together to form protons and neutrons. They all interact via the strong force and, because the strong force binds them up so tightly, they are never found on their own but only inside of composite particles like protons and neutrons. There are six types of quarks, although only two of them (up and down) show up in protons and neutrons. They are shown in purple on the table. The quarks also carry electric charge, which is why the proton has a net positive charge.

2. **The Electron and Its Siblings:** As far as we know, the electron is a fundamental particle. It does not interact via the strong force but it is electrically charged, which is why it is attracted to protons and ends up bound up inside of atoms. It is also has a small mass compared to protons and can in some cases float easily from atom to atom, giving rise to phenomena such as electrical currents. There are two siblings of the electron known as the muon and the tau: they are heavier than the electron but otherwise have identical properties. Together, these particles are known as the **charged leptons**, and they form the top green row in the table.

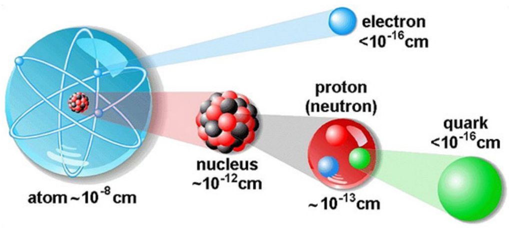

Figure 6: Illustration of how the atom is made of protons, neutrons, and electrons; the protons and neutrons are in turn made up of various types of quarks.

3. **The Neutrinos:** These are bizarre particles that have extremely tiny masses (about a million times smaller than the electron mass) and only interact via the weak force. Because of this, they are incredibly hard to detect and their properties are consequently less well understood than other Standard Model particles. The Nobel Prize in 2015 was awarded to the experimental groups that discovered that the neutrinos are not massless but have a tiny, non-zero mass; this is in contradiction with the originally posited version of the Standard Model, which predicts that neutrinos are exactly massless. Neutrinos are also known as the **neutral leptons**, and they form the bottom green row in the table. There are three types of neutrinos, one corresponding to each of the types of charged leptons.

4. **The Force Carriers:** In a quantum mechanical theory, the energy transmitted between objects is passed in little lumps. This may be familiar to you if you have studied the photoelectric effect or the double-slit experiment for light. When two particles interact via a force, that force is transmitted by a third particle known as the "force carrier" or, more typically among physicists, the "mediator". The technical name of such a particle is a **gauge boson**, so you might see that name used too. The mediator of the EM force is the **photon**, the mediator of the strong force is the **gluon**, the mediator of the gravitational force is the **graviton**, and the mediators of the weak force are the W **and** Z **bosons**. These make up the red column of the table; note that gravity is not actually included in the Standard Model because of its irrelevance for elementary particle interactions, and so the graviton is absent in Fig. 7.

5. **The Higgs Boson:** This particle was first predicted in the early 1960s and was sought experimentally all the way until its discovery in 2012 (and the subsequent award of the Nobel prize in 2013). Given the way that SM particles interact with the known forces, without a Higgs boson it would be impossible for any of the other Standard Model particles to have a mass. Understanding the details of the Higgs boson and its action in the Standard Model is beyond the scope of these notes, but it is safe to say that it is the linchpin without which our whole understanding of elementary particles would fail. This is the yellow final column of the table. Its label in the table as a "scalar" boson means that Higgs bosons have no spin, unlike the force carriers or gauge bosons which do have spin.

Together with the forces, this small set of particles is sufficient to explain the formation of atoms and, consequently, all of the phenomena we see in every other branch of physics.

While the idea of distilling everything we know in physics down to a handful of forces and particles (and their corresponding force equations) is extremely appealing, the goal of particle physics is not simply to construct a recipe book of known particles and interactions. Rather, we want to use our knowledge of fundamental particles to answer deep questions about *why* the world works the way

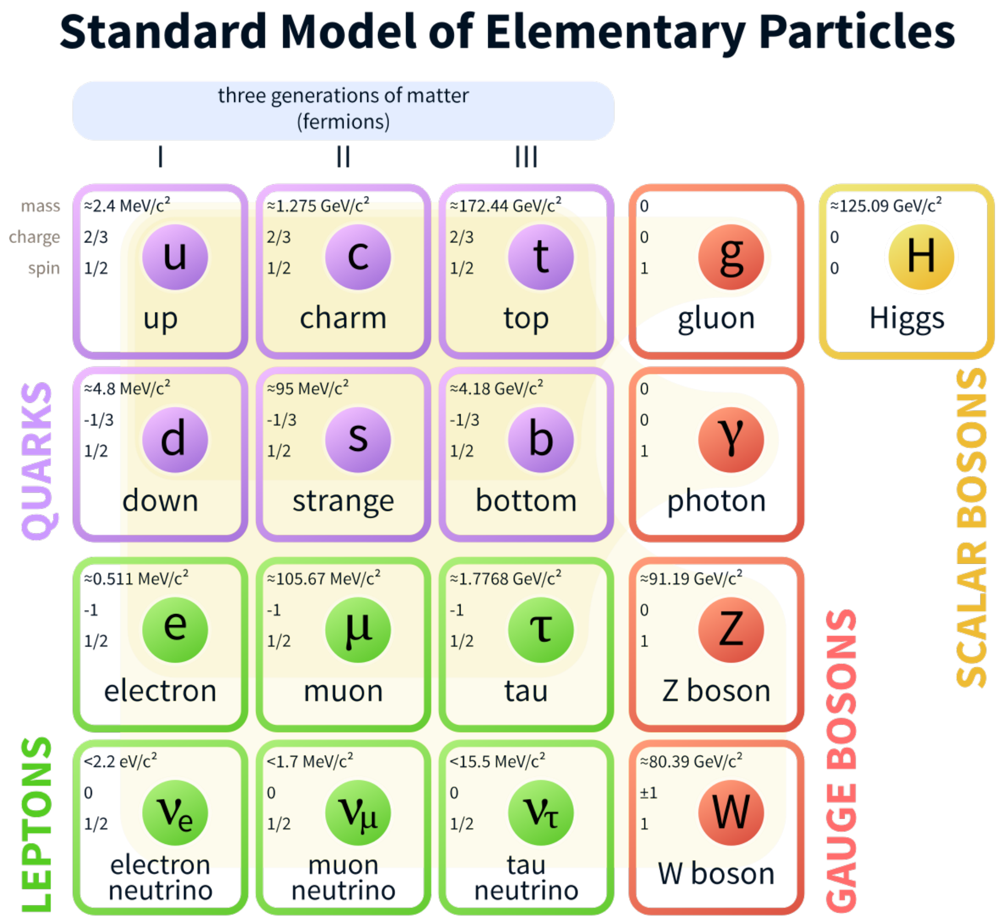

Figure 7: The particles of the Standard Model of particle physics and their properties. The name and symbol of each particle is found in the centre. At the top left is the mass of the particle, followed by the electric charge (in units of the electron charge magnitude), and the spin (in units of Planck's constant). Spin is the intrinsic angular momentum of a fundamental particle (image source: Wikipedia).

it does. This could, for example, lead us to deeper insight into the nature of our world and our understanding of the most basic principles of physics. Right now, we have a satisfactory description of most physical phenomena based on the particles and forces in the SM. However, the SM seems much more complicated than it needs to be: for example, we have no understanding of why there are so many forces, whether they originate in some more fundamental unifying force, and why certain particles are charged under one type of force but not another. We also don't understand why there are so many different particles: for example, the matter particles in the SM appear to be grouped into three "families" or "generations" corresponding to each column, but we don't understand why that is. Nor do many of these heavier generations seem to play much role in physics as it unfolds in the Universe today.

If we are optimistic, we can imagine that there is more to learn about particle physics by considering an analogy with the discovery of atoms and atomic structure. In the early 19th Century, it was realized that atoms are the building blocks that make up all chemicals and compounds known in chemistry. However, there are a relatively large number of such building blocks (currently 118 and counting). It is hard to make sense of such a large number of them, but once Mendeleev organized them into the Periodic Table, it was possible to notice that atoms in a given column shared certain properties: the inert Noble gases all fit together, while the highly reactive alkali metals share another column. We now understand that this structure has to do with the number of valence electrons in the atom: once we understand the sub-atomic particles that comprise the atom, we can understand and make predictions about how atoms made of different numbers of electrons, protons, and neutrons behave. It is clear that the organizational structure of the elements reflects some deeper truths about how the atoms themselves are configured. Perhaps the assortment of particles and forces we see in Fig. 7 is a hint of a deeper structure that we are as of yet unaware! Certainly, we can be heartened by the fact that atoms were once thought of as fundamental, and so we may yet discover hidden secrets inside of quarks, leptons, or force mediators.

There are other reasons to believe that there is more to the story than the Standard Model. Astrophysical observations tell us that there seems to be a lot of **dark matter** in the universe, which does not give off any light but whose gravitational interactions influence the dynamics of stars and other visible objects. In fact, there is over five times more dark matter in the Universe than "regular" matter made of atoms! As far as we can tell, none of the known SM particles fit the bill for explaining dark matter: they either interact too strongly with other types of SM particles, decay away too quickly after the Big Bang, or are too low in mass to properly explain the dark matter structures we see in the Universe. It is very likely that dark matter is some *new* kind of particle that we currently know nothing about[4]. Because the gravitational force is universal (meaning it cares only about the total mass/energy in an object, not the types of particles that make up that mass/energy), we don't currently know much about dark matter beyond how much mass-energy there is in dark matter in total. It may be that dark matter is a new particle that only interacts with us via the weak force (which could explain why we haven't seen it yet), or it may not interact with us by any of the known forces. Indeed, it is hoped that our observation of dark matter will open a window to a whole new world to be explored.

I also want to mention another mystery that is near and dear to me: the puzzle of why the universe seems to be made only of matter and not antimatter. So far, I haven't made any mention of antimatter, but every particle has a corresponding antiparticle which has the same mass but all of its charges have flipped sign: for example, the electron has a corresponding **positron** (or anti-electron), which has electric charge $+e$ instead of $-e$. When a particle and antiparticle meet, they annihilate into energy or other forms of matter. If the Universe were made of equal parts of matter and antimatter, all the matter and antimatter would have annihilated away, leaving nothing but a bath of glowing light: no atoms, no stars, no galaxies, no us. Our equations for the four forces treat matter and antimatter on nearly the same footing, and so the SM cannot explain why the universe we see is populated by protons, neutrons, and electrons instead of antiprotons, antineutrons, and positrons, or how we ended up with that excess of matter over antimatter in the first place. The matter-antimatter asymmetry is another hint for the existence of new particles and forces that treat matter and antimatter differently.

What we see from these last two examples is that, even though particle physics is the study of very tiny objects (namely, elementary particles), we can use it to understand much larger objects

[4]Because we detect dark matter through its gravitational pull on regular, visible matter, it is also possible that the observations we see are *not* actually signs of some invisible matter but actually a signal that our theory of gravity is not correct. An alternative to dark matter theories involve modifying gravity. This is typically difficult to do in a way that matches all of the various cosmological observables we see at different moments after the Big Bang, but it's not completely excluded as an alternative explanation for dark matter. Most physicists, however, accept the dark matter theory.

such as galaxies and beyond. The reason is that the early universe was a much more compressed, hotter, and denser place, and all of the particles in the SM were being produced, destroyed, and participating in rapid collisions. All of these collisions influenced the evolution of the universe as a whole and everything in it, and leave definitive imprints in various features we observe today. There is a very clear link between the tiniest objects in the known universe and its large-scale structure, and so particle physics allows us to study very deep and profound questions about our world in a unified framework.

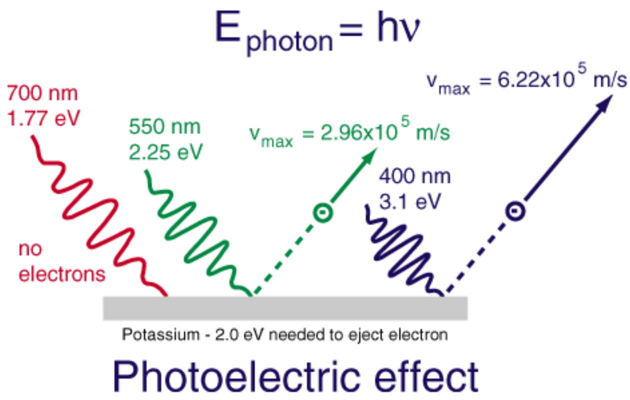

Figure 8: The photoelectric effect: the electromagnetic field is made up of lumps of mediator particles called photons. When light is shone on a metal, individual electrons in the metal only interact with a single photon at a time. Unless the photon's energy is sufficiently large to eject the electron from its host atom, no electrons are removed from the metal. The photon energy depends on the frequency, or equivalently, the colour of light (image source: Georgia State University Hyperphysics).

2 The Three Forces of the Standard Model

In this section, we discuss each of the forces in more detail. We focus on the three most important forces for particle physics, namely the electromagnetic, weak, and strong forces; in strict terms, the Standard Model (SM) encompasses only these three forces. The gravitational force is typically irrelevant for the study of particle physics except in its classical formulation as applied to large objects; however, we briefly discuss it for the sake of completeness in Sec. 2.4.

Occasionally, we will use results from elementary quantum theory or special relativity. If these concepts are new to you (or rusty), there is a brief discussion in Appendix A; you may also want to talk to your research mentor or look around online for more resources.

2.1 Electromagnetism

Electromagnetism is the most familiar of the forces in the SM and so it is natural to start here. There are a few things that we know already about electromagnetism. This force arises from the dynamics of electric and magnetic fields, which are in turn sourced by the presence and motion of electrically charged particles. Particles can posess positive or negative electric charges (or be neutral, equivalent to zero charge), and the electromagnetic force can be either attractive between oppositely charged particles, or repulsive between similarly charged particles.

The classical electric and magnetic fields are composed of large numbers of individual lumps of energy known as **photons**. The particulate/quantum nature of the electromagnetic field was first understood by reference to the photoelectric effect shown in Fig. 8. The photoelectric effect is the observation that, when electromagnetic waves (*i.e.,* light waves) hit a metal plate, the ability of the light to eject electrons from the metal depends on the *frequency* of the light, not the intensity of the light. The reason is that the beam of light is made up of many photons, each of which has an energy that depends on the frequency of light. If the energy in a single photon is insufficient to eject an electron from an atom, then no electric current is emitted from the plate.

The photon itself is massless; according to special relativity, only massless particles travel at the universal speed of light c (in every reference frame), and since light itself is made of a collection of

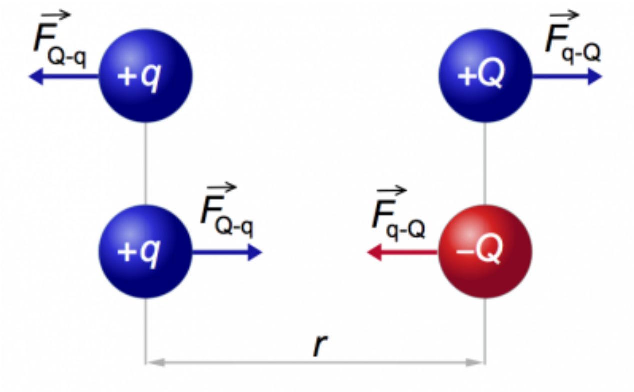

Figure 9: The coulomb force acting between like and opposite-charged particles (image source: Wikipedia).

photons, it is understood that photons are massless too[5]. Our symbol for the photon is the lower-case Greek letter gamma, γ, and for this reason high-energy photons are sometimes referred to as gamma rays. The photon is electrically neutral and is its own antiparticle.

Apart from specifying the mass of the mediator, we also need to quantify the strength of its interaction with elementary charged particles. This cannot be derived from the theory and needs to be measured experimentally. Recall that the electron has charge $-e \approx -1.6 \times 10^{-19}$ C (in SI units). Suppose we have two such electrons separated by a distance r; then, we know that a repulsive force exists between them as shown in Fig. 9. For concreteness, we fix the left-hand charge at the origin and compute the electrostatic force $\vec{F}(\vec{r})$ on the right-hand charge in the diagram which has position $\vec{r}$. Because the force is conservative (*i.e.*, energy is conserved), we can write the force as the gradient of a potential, $\vec{F}(\vec{r}) = -\vec{\nabla} U(\vec{r})$ (or if you haven't taken multivariable calculus, recall that the work done in 1D on an object by a force is $\Delta W = \int dx' \, F_x(x')$. We can define the change in potential energy as $\Delta U = -\Delta W$, since positive work being done on an object reduces its potential energy and increases its kinetic energy. Consequently, the force can be written as $F_x = -dU/dx$ where U is potential energy).

The Coulomb force between two electrically charged objects is

$$\vec{F}(\vec{r}) \quad = \quad \frac{e^2}{4\pi\epsilon_0} \frac{1}{r^2} \hat{r}, \tag{2}$$

where r is the magnitude of the position vector, $\hat{r}$ is the unit vector pointing in the radial direction, and $\epsilon_0 \approx 8.85 \times 10^{-12}$ C^2/m/J is the permittivity of free space. Integrating the force along a radial path, we obtain the result for the potential energy between two electrons:

$$U(\vec{r}) \quad = \quad \frac{e^2}{4\pi\epsilon_0} \frac{1}{r}. \tag{3}$$

This is the usual result that the electrostatic repulsion becomes stronger as the charges are brought closer together, and hence more energy is needed to bring the electrons close together.

We can define a *dimensionless* physical quantity known as the **fine structure constant**, α:

$$\alpha \quad \equiv \quad \frac{e^2}{4\pi\epsilon_0 \hbar c}, \tag{4}$$

[5]The universality of the speed of light has only been tested to finite precision, and so all we can say for sure is that the photon has a mass $< 10^{-27}$ times the proton mass, but that is very small!

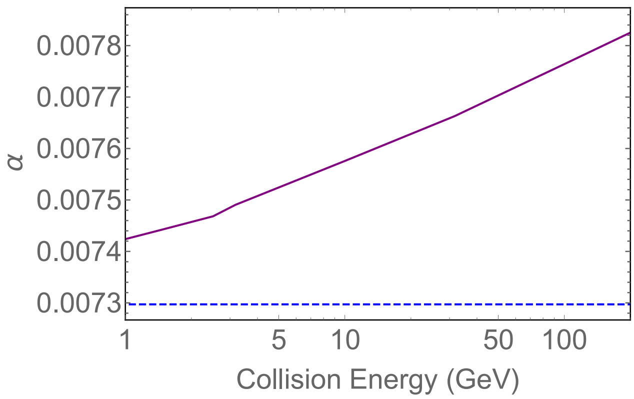

Figure 10: The solid purple curve shows the value of the electromagnetic fine structure constant α measured at different particle collision energies. The blue dashed curve shows the zero-energy value of $\alpha(E=0) \approx 1/137$. It is apparent that α grows logarithmically with energy.

where $\hbar = h/2\pi \approx 1.05 \times 10^{-34}$ J $\cdot$ s is the reduced Planck's constant and $c \approx 3 \times 10^8$ m/s is the speed of light[6]. We see that the electrostatic potential is proportional to α,

$$U(r) = \frac{\hbar c \alpha}{r}. \tag{5}$$

In this form, the potential is specified entirely in terms of the *fundamental, universal* constants c and $\hbar$, plus a single dimensionless number that characterizes the strength of the force. Indeed, we can speculate that any force X gives a potential between "charged" particles of the form

$$U_X(r) = \frac{\hbar c \alpha_X}{r} \tag{6}$$

where α_X is a dimensionless quantity characterizing the strength of force X If α_X is small relative to 1, then the force is a "weaker" force, whereas if $\alpha_X > 1$ then the force is "strong". Indeed, the Coloumb form of the potential is really only valid for $\alpha_X \lesssim 1$.

Returning to electromagnetism, we have

$$\alpha \approx 7.297 \times 10^{-3} \approx \frac{1}{137}. \tag{7}$$

The latter fractional approximation is a useful way of remembering the value of α and is used frequently in approximations of its value. α is one of the most precisely measured physical constants and is known to better precision than one part in one billion!

The mass of the photon and α are basically all we need to do calculations involving the electromagnetic force. We will explore in more detail the elementary interactions of photons and electrically charged particles in Secs. 3 and 4. For now, however, there is one additional subtlety that is necessary for understanding the behaviour of the forces: in spite of the fact that we call α a "constant", it is not *actually* a constant but rather it depends on energy. If an interaction between a charged particle and a photon occurs at an energy scale E_{int}, then we actually have to evaluate α at that scale $\alpha(E_{\text{int}})$. We call this phenomenon a **running interaction strength**, or running coupling.

[6] Generally, the appearance of Planck's constant in a formula suggests that it is a quantity relevant to quantum mechanics, while the appearances of the speed of light suggests that special relativity is important. Indeed, particle physics is studied in the regime where both quantum mechanics and special relativity are needed to get the right answer because we look at very small and very fast objects (like in particle physics).

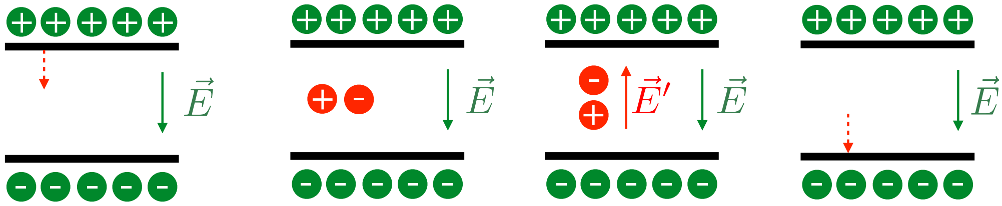

Figure 11: Charge screening that leads to an energy dependence of α. In the left pane, we have a configuration of electric charges that leads to an electric field $\vec{E}$ pointing down. A photon (red dashed line) can be emitted from one plate and absorbed by the other, mediating the charge. In the second pane, the photon very briefly splits into an electron-antielectron (or positron) pair. In the third pane, the electron and positron pair drift according to the electric field, setting up an opposing electric field $\vec{E}'$ that tends to cancel (or **screen**) the initial field $\vec{E}$. In the fourth pane, the electron and positron re-annihilate into the photon, which continues on its way. Because this electron-positron pair appearance and disappearance occur rapidly all over the interior between the plates, this leads to a continual partial cancellation/screening of $\vec{E}$. Low-energy photons have large wavelengths, and so are sensitive to this screening effect, whereas high-energy photons have short wavelengths and can only see a tiny part of the volume between the plates at a given time, reducing the screening effect. This is why the electromagnetic force is stronger over higher energies/shorter distances.

For example, when electromagnetic interactions were measured at the energy scale of the former LEP collider at CERN, $E_{\text{LEP}} = 91$ GeV[7], it was found that $\alpha(E_{\text{LEP}}) \approx 1/129$ instead of the value $\alpha(E = 0) \approx 1/137$. Therefore, we find that the strength of the electromagnetic interactions *grows* at higher energies! We show a plot of the running value of α in Fig. 10. Confusingly, we often still refer to the energy-dependent value of α as the fine-structure "constant", but because this wording is deeply embedded in the particle physics vocabulary, it is something you may see frequently.

Exercise: (a) What is the energy of a photon of blue light ($\lambda = 475$ nm)? Express your answer in J and eV. (b) What is the wavelength of light associated with a photon of 1 keV in energy? What is the name for this part of the electromagnetic spectrum? (c) At the LHC, particles are collided with an energy of 7 TeV. What is the wavelength of a 7 TeV photon?

(Optional, for those who have studied some electromagnetism) The fact that a *dimensionless* quantity that characterizes the strength of the interaction has a secret dependence on energy may be a bit strange: since α has no dimension of energy (or anything else, for that matter), we do not expect the interaction to be energy dependent. The reason for the energy dependence is as follows: if a photon is being exchanged in an interaction between two particles, there is small chance that the photon will briefly split into an electron-positron pair before re-combining and continuing on its journey. This electron-positron pair can act like an electric dipole, aligning its charges so as to cancel out the electric field in the intermediate region. This gives rise to a **screening** of electric charges, such that at large distances electric charges seem smaller than they otherwise would; see Fig. 11. In quantum mechanics, longer distances and wavelengths are associated with smaller energies (see Appendix A.1.1), and so the effective strength of EM interactions is suppressed at lower energies[8].

[7] In particle physics, we often use units of energy in **electron-volts**, or eV. The conversion to SI units is 1 eV $\approx 1.6 \times 10^{-19}$ J; it is the energy that one electron acquires if it crosses one Volt of potential difference. Therefore, 91 GeV $\approx (91 \times 10^9$ eV$) \times (1.6 \times 10^{-19}$ J/eV$)$.

[8] This is precisely the same reason that an electric field inside of a dielectric, such as a capacitor, is smaller than it

At higher energies, you are probing smaller volumes around each charge and so this screening is less effective, resulting in a larger value of α.

2.2 The Weak Force

The weak force is, well, weak, and so it does not have many observable effects in our day-to-day life. Its only manifestation at low energies is that it is responsible for certain radioactive decays such as beta decay, where a neutron decays to a proton plus an electron and anti-neutrino ($n \to p^+ e^- \bar{\nu}_e$). Just like the electromagnetic force, the weak force is specified by two quantities: the mass of its force carriers, and the dimensionless coupling "constant" that characterizes the strength of the interaction. For the weak interaction, we define a **weak structure constant**, α_{w}, which is analogous to α for the EM force.

The force carriers of the weak force are called the **$W^\pm$ and Z bosons** (see Fig. 7). The $W^\pm$ bosons are electrically charged and are antiparticles of one another (*i.e.,* the antiparticle of the W^+ boson is the W^- and vice-versa), while the Z boson is electrically neutral and is its own antiparticle. Unlike the EM force, however, the weak bosons are *not* massless. In fact, they are very heavy. In particle physics, we often use units of mass of GeV/c^2. The proton has a mass $M_p = 0.938$ GeV/c^2, and so masses in GeV/c^2 are approximately the same as multiples of the proton mass. The weak boson masses are

$$M_{W^\pm} = 80.4 \text{ GeV}/c^2, \quad (8)$$
$$M_Z = 91.2 \text{ GeV}/c^2. \quad (9)$$

It requires an enormous amount of energy ($E = M_W c^2 = 80.4$ GeV, equivalent to electrons accelerated with 80 billion volts) to physically produce these particles, and they were not discovered until a sufficiently energetic proton collider experiment was built at CERN in the 1980s.

Exercise: Show that units of GeV/c^2 *have the correct dimensions for mass.*

Exericse: A wall outlet in North America has an electric potential difference of 120 V between the prongs. How many times larger of a potential difference is needed to accelerate an electron so that its energy is equal to the rest mass energy of a proton?

The dimensionless coupling constant as measured by the LEP experiment is

$$\alpha_{\mathrm{w}}(E_{\mathrm{LEP}}) \approx 0.03. \quad (10)$$

Note that $\alpha_{\mathrm{w}} > \alpha$ from electromagnetism! This is confusing: why do we call the weak interactions "weak" when they are actually *stronger* than the EM interactions?

The resolution to this puzzle lies in the fact that the $W^\pm$ and Z are very heavy and consequently have a large rest energy, $E_W = M_W c^2$ (and similar for the Z boson), whereas the photon is massless. In quantum mechanics, objects of high energy are smaller in size too (see Appendix A.1.1). If particles are colliding at energies well below $M_W c^2$, then there is not enough energy to physically produce a $W^\pm$ boson that moves a finite distance; instead, we rely on the $W^\pm$ and Z bosons temporarily flitting into existence over a very tiny distance to mediate the force. Thus, the weak force mediated by the $W^\pm$ and Z bosons will extend only over a very short distance. This is in contrast to a photon, which

would be in empty space. The fact that photons can split into particle-antiparticle pairs of charged particles in vacuum means that empty space itself has some dielectric properties.

has no mass and therefore no minimum energy; the photons can be produced at will and travel over great distances, leading to a *long-range* force.

We can use Planck's constant $\hbar$ and the speed of light c to convert the mass, M, into a quantity with dimensions of distance,

$$R \approx \frac{\hbar}{Mc}. \tag{11}$$

This is the typical spatial extent of the force carrier particle (the $W^{\pm}$ or Z), and so the weak force can only act over at most a distance R. This is due to the connection in quantum mechanics between objects with large energies and their small associated sizes.

Exercise: The mass of the W boson is 81 GeV/c^2. *Compute the corresponding distance R_W in metres.*

It is evident that the weak force is only relevant over very tiny distances! The Coulomb form of the potential energy, $\sim 1/r$, does not apply for the weak force over long distances, $r \gg R$, since the $W^{\pm}$ and Z bosons only act on distances $r \lesssim R$. It turns out that the correct form of the potential energy is instead

$$U(\vec{r}) = \frac{\hbar c \alpha_{\rm w}}{r} e^{-r/R} = \frac{\hbar c \alpha_{\rm w}}{r} e^{-Mcr/\hbar}. \tag{12}$$

This form of the potential energy is known as the **Yukawa potential**. For $r \gg R$, the weak force is totally irrelevant in attracting or repelling two objects because of the exponential factor, and this explains why we don't see any macroscopic effects of the weak force. However, the overall strength of the force is still determined by the coupling constant $\alpha_{\rm w}$.

Does the weak force have any effect at all on low energy particles? The answer is yes, because the weak force does lead to the radioactive beta decay of nuclei. To understand this, though, we have to think a bit more about the connection between energies and distance scales. Let us suppose that we have two particles, A and B, with equal energies $E = E_A = E_B$ travelling in opposite directions (see Fig. 12). Quantum mechanics tells us that there is a characteristic size associated with each object, $r_A, r_B \sim \hbar c/E$ (again, see Appendix A.1.1). If each particle's energy is much less than the mass-energy in the W and Z bosons, $E \ll M_W c^2$, then we have that $r_A, r_B \gg R_W$. In other words, the spatial extent of each of particle A and B is larger than the range of action of the weak force.

How do we interpret this large "cloud" or "blob" corresponding to each particle? If you have studied chemistry, you may know about the *orbital* interpretation of the electron: the electron occupies a large volume but the electron itself is pointlike, with the orbital specifying the probability that you will see the electron at any particular point. Similarly, we can think of the spatial profile of the particles A and B as related to the regions where we are likely to find particle A or B. If, as they cross, the particles happen to lie within the same tiny volume of radius R_W, then they *can* interact via the weak force. Therefore, the probability of A and B interacting via the weak force is proportional to the probability that they occupy the same tiny volume of space, relative to the overall spatial extent of the particle.

The total spatial extent of the particle follows from $r_A, r_B \propto 1/E$, whereas the spatial range of the weak force mediated by the W boson is $R_W \propto 1/(M_W c^2)$. We can therefore see that the probability of interacting via the weak force, as a function of the energies of the colliding particles, goes like

$$\mathcal{P}_{\rm weak}(E) \propto \left(\frac{R_W}{r}\right)^n \propto \left(\frac{E}{M_W c^2}\right)^n, \tag{13}$$

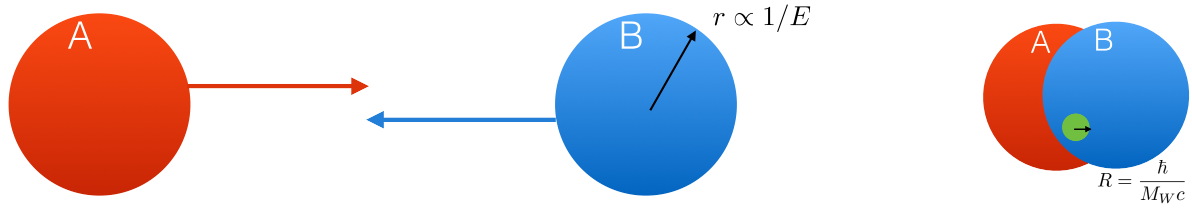

Figure 12: Illustration of a collision mediated by the weak force. Particles A and B, each with energy E, are headed towards one another (left pane). Quantum mechanics says that particles with energy E are actually distributed over a volume of radius $r \sim \hbar c/E$, and there is some probability of finding the particle at any point in the shaded region. However, for the weak force to act, both of particles A and B have to be located within a much smaller distance $R = \hbar/(M_W c)$ indicated by the green blob (right pane). Because the probability of finding both A and B within a distance R is small, this leads to a suppressed probability of interaction, prob $\propto (R/r)^n \propto (E/M_W c^2)^n$ for some positive power n.

where $n > 0$ is some positive power. We see that the *lower* the energies of the particles in the collision, the more spatially extended the objects are and the *smaller* the probability that they will happen to be in the same tiny volume of radius R and can interact via the weak force.

We can now explain why the weak force is known as weak even though the dimensionless number characterizing its strength is larger than for electromagnetism. The reason is that, for energies much smaller than the mass-energy of the force carriers, the colliding particles are so spread out as to make the probability of both particles being within a distance R_W of one another vanishingly small. As a result, the force becomes very feeble as $E \to 0$ with only a tiny probability of the force doing anything to a physical system. This is why radioactive beta decay occurs relatively slowly: the energies released in the decay are much smaller than the mass-energy of the W and Z bosons.

(Optional) Notice that something funny happens for $E \gg M_{\text{weak}} c^2$: the probabilities seem to grow without limit! This is very bad: after all, the maximum total probability of something happening should be 1, whereas Eq. (13) suggests that the probability goes to infinity for large energy[9]! This is very bad and a sure sign that something is wrong with the calculation. However, if $E \gtrsim M_W c^2$, then we can actually start producing the physical $W^{\pm}$ and Z bosons in our collision. In other words, particles A and B in the collision become small enough that the weak force can act over the entirety of the particles. Therefore, in this limit the weak force is no longer pointlike and can act over a distance (just like the EM force), and the force simply reverts to a Coulomb-like $1/r$ potential.

Finally, we remark that α_{w} has a dependence on the energy of interaction just like the EM force, although the effect is the opposite: α_{w} gets smaller at larger energies. Its evolution with energy is pretty minor, however, just like for α_{EM}. At high-energy colliders such as the Large Hadron Collider (LHC) with collision energies $E \gg M_W c^2$, the weak force is actually more important than the EM force in spite of its name. In Sec. 5.3, we will explore more of the specific processes mediated by the weak interactions.

2.3 The Strong Force

Finally, we turn to the strong force. It is responsible for holding quarks together inside of the protons and neutrons (collectively referred to as nucleons), and also for nucleons to stick together inside of

[9]Because the conservation of probability in quantum mechanics is associated with the states in the theory evolving via unitary transformations, such behaviour is often called a breakdown of unitarity.

the nucleus (for a reminder, see Fig. 4). However, we don't see any evidence for the strong force on distances larger than the nucleus ($\gtrsim 1 \times 10^{-15}$ m). This is similar to how we don't see any evidence for the weak force at long distances. According to our earlier discussion, it is therefore tempting to associate the strong force with a strong-force mediator of mass $M_\mathrm{s} = \hbar/(\mathrm{fm}\, c) \approx 0.2$ GeV/c^2 by analogy with the W and Z bosons. However, in experiments we find no such simple force carrier; instead, around this mass we see a HUGE number of new particles appear! This puzzled physicists through much of the 1950s and 1960s.

We now know that the force carrier of the strong force, known as the **gluon**, is actually *massless* like the photon. This seems very strange: we certainly don't see a long-range force associated with the gluon like we see for the photon (recall that massless force carriers can act over arbitrarily long distances!). We will soon find the resolution to this paradox, but first let us finish describing what we now understand about the strong force. Particles that interact via the strong force carry charges much like the EM force, except that there are *three* types of charges. They are known as **colour charges**, and so the theory of the strong force is known as **quantum chromodynamics (QCD)**, namely the dynamics of colour[10]. The three types of colour charges are commonly referred to as green, red, and blue. The antiparticles carry **anti-colour**, namely anti-red, anti-blue, and anti-green.

In electromagnetism, we know that a charge $+Q$ and a charge $-Q$ combine to give a net neutral state. Similarly, one colour is neutralized by its corresponding anti-colour. For example, red combines with anti-red to give an colour-free state. There is another combination that gives a neutral charge under the strong force: a combination of red + green + blue *also* gives a neutral overall charge, and the same is true for anti-red + anti-green + anti-blue. In other words, the combination of every colour is equivalent to no colour charge. This is not possible in electromagnetism because there is only one "type" of charge that comes in two signs.

With electric charges, it is easy to figure out which combinations attract (opposite-sign charges) and repel (same-sign charges). With the strong force, it is much more complicated because there are several different colours. It is made even more complicated by the fact that the gluons themselves carry a colour charge. So, when two particles carrying colour charge exchange a gluon, their colours change after emitting/absorbing the gluon. Very confusing indeed!

However, we will leave these subtle points aside and define the force according to the mass of the mediator and a dimensionless coupling strength just like before. I have already said that the gluon is massless (although we have yet to explain why we don't see a long-range strong force). We can write down a potential energy mediated by the force (up to a sign that signals attraction vs. repulsion, which we leave out):

$$U_\mathrm{strong}(\vec{r}) \quad = \quad \frac{\hbar c \alpha_\mathrm{s}}{r}, \tag{14}$$

where α_s is the **strong structure constant**. It is, once again, dimensionless:

$$\alpha_\mathrm{s}(E_\mathrm{LEP}) \quad \approx \quad 0.1, \tag{15}$$

and, true to its name, it is *larger* than the corresponding interaction constants for the other forces, $\alpha_\mathrm{s} > \alpha_\mathrm{w} > \alpha$.

A crucial difference between the strong force and the EM force is in its running, or how it changes as a function of collision energy. Like the weak coupling, α_s gets *smaller* at higher energies,

[10] Of course, there are not colours in the conventional sense; they are simply "cute" names for keeping track of the three types of charges. You'll soon find that the world of particle physics is full of weird names that people made up, and some choices are more helpful or problematic than others. If you are interested in learning more on the connections between particle physics, astrophysics, and the social dynamics of our field and broader society, you should check out *The Disordered Cosmos* by Chanda Prescod-Weinstein.

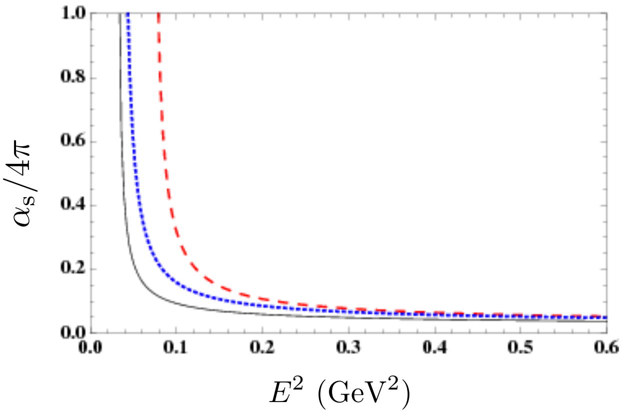

Figure 13: The running of the strong interaction coupling constant, $\alpha_s(E)$, as a function of energy. The strong force structure constant gets *smaller* at larger energies, unlike in electromagnetism (see Fig. 10). It also grows stronger at lower energies, blowing up at $E \sim 0.2$ GeV (corresponding to the force becoming extremely strong). The different curves correspond to different theoretical calculations (image source: `arXiv:1208.5636`).

and consequently the strong interaction gets weaker at higher energies (see Fig. 13). This property of the strong force is known as **asymptotic freedom**, because at very high energies the strong force essentially turns off and stops acting on anything. The reason for this different behaviour relative to the EM force is that the gluons themselves carry colour charges (unlike the photon), and this together with the higher multiplicity of colour charges causes its interaction strength to diminish at higher energies.

The flip side of the asymptotic freedom coin is that the strong interactions get *stronger* at lower energies. In fact, there is a particular energy scale called $\Lambda_{\rm QCD}$ at which α_s becomes formally infinite! You can see the curve going to infinity in Fig. 13. Of course, there is not actually a physical quantity that becomes infinite; what this is telling us is that the force becomes so strong that our calculation of the running of $\alpha_s(E)$ breaks down and starts giving unreliable answers. The numerical value is $\Lambda_{\rm QCD} \approx 0.2$ GeV, a number that should be familiar to us from earlier in the section as the energy at which we start seeing massive numbers of particles appearing.

What happens for $E \lesssim \Lambda_{\rm QCD}$ (and, consequently, longer distances) is that the strong force is so strong that all particles that carry colour charge are immediately attracted to one another and bound up into states that are *colour neutral*. One way we can imagine this is to think that the Coulomb-like potential in Eq. (14) gets replaced on long distances with a potential[11] with the property that $U(\vec{r}) \to \infty$ for $|\vec{r}| \to \infty$. In other words, if you try to "infinitely" separate two objects with colour charge, an inifinite energy is required! This is in contrast with the Coulomb potential for electromagnetic or gravitational forces, where $U \to 0$ for $|\vec{r}| \to \infty$. This property of the strong force is known as **confinement**. It has never been rigorously proven, but there is very strong experimental evidence in its favour[12].

Confinement is why we never see any free, individual particles carrying colour charge (such as quarks and gluons): they are all strongly bound to *other* quarks and gluons to form colour-neutral bound-state particles. This also explains why we don't see a massless gluon, and hence no long-range

[11]You may wonder why this is the case. As a rough argument, it is fair to say that Eq. (14) is the first term in a series expansion in the small parameter α_s. Once we have $\alpha_s \gg 1$, the series expansion is no longer valid and so the potential has a very, very different form.

[12]If you can rigorously prove it, though, then the Clay Mathematics Institute will award you a million dollars!

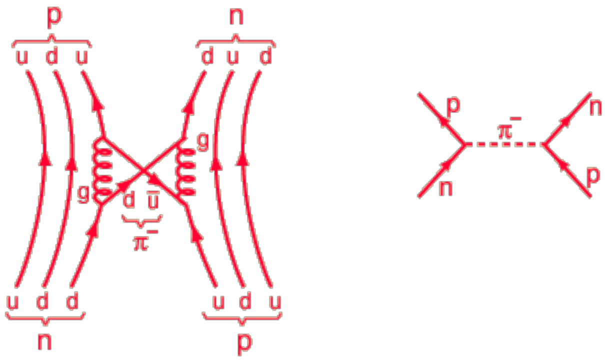

Figure 14: The residual strong force originates from gluons within the proton or neutron forming intermediate hadrons (left pane), known as pions. These pions mediate the force from one proton or neutron to another (right pane). The gluon cannot directly pass from one proton or neutron to another due to the confining force which keeps colour-charged states inside of individual protons or neutrons (image source: Georgia State University HyperPhysics).

force: the gluons are also bound into colour-neutral particles. The colour-neutral bound states that are formed after confinement are called **hadrons**, and the most familiar examples are the proton and neutron. There are others, for example, such as pions and kaons. We will discuss these further in Sec. 3.4.

Exercise: In introductory mechanics, we often discuss the potential energy due to a stiff spring, $U(x) = kx^2/2$. Does this potential exhibit confinement? In terms of a physical spring, what would happen if we put a very large amount of energy ($U \to \infty$) into the system?

(Optional) It turns out that confinement is not exactly the end of the story. In some sense, the strong force is sufficiently strong that it "leaks" out of the individual hadrons. This is known as the **residual strong force** and it is responsible for having the keeping the protons and neutrons stuck together inside of the nucleus. Suppose we have two protons close together. Naïvely, we would expect them to repel because they are both positively charged. However, gluons in one proton (in some sense) try to escape the proton to interact with the quarks in the other proton; because the gluons carry colour charge, they cannot exist outside of a hadron but they can form into a *different* hadron, which is emitted by one proton and absorbed by the other. The process is illustrated in Fig. 14. In the residual strong force, some of the hadrons act as mediators of the strong force between two nucleons. The most important of these are the pions, which have masses $M_\pi \approx 0.14$ GeV$/c^2$, and these explain the finite range of even the residual strong force in attracting two nuclei.

2.4 (Optional) What About the Fourth Force?

In the introduction, we talked about four fundamental forces: electromagnetism, the strong force, the weak forces, and gravity. One reason for including gravity in such a list is that it is a force familiar to everyone (whether a physics student or otherwise), and so its exclusion should naturally invite questions. However, gravity is a very *feeble* force. One way to see this is to look at a bar magnet that is picking up a set of paper clips: the force of the tiny bar magnet is sufficient to overcome the gravitational pull of the *entire Earth* and lift the paper clips up off a table! When we look at particles at the subatomic level, they are extremely tiny and light (by the standards of planets and stars), and so the force of gravity is utterly negligible. For these reasons, gravity is typically not explicitly included in the list of forces of the Standard Model.

Still, there is a lot of interest in how gravity applies to elementary particle physics. Indeed, you may have read in a popular science book that "gravity and quantum mechanics are incompatible", and that if we try to apply the rules of quantum mechanics to the physics of gravity, we end up with totally nonsensical answers. This may be a good story to tell when selling books, but unfortunately it is inaccurate. There is actually nothing inherently incompatible about gravity and quantum mechanics, and so we may include it in the Standard Model if we wish. We can even compute quantum mechanical corrections to the classical gravitational force, and they are perfectly well defined (if utterly negligible as far as experiments are concerned).

If we do include gravity as a fundamental force, then we must consider its force carrier particle. The mediator of the gravitational force is called the **graviton**, and it is massless. This is again consistent with our understanding of gravity as a force with an arbitrarily long range, $U(\vec{r}) \propto 1/r$. Unlike photons, gluons, and $W^{\pm}$ and Z bosons, we have never actually seen a graviton, and this is because gravity interacts so very weakly with elementary particles (due to their tiny masses) compared with other forces. Nevertheless, there is no good reason to believe that gravitons do not exist given the analogy with each of the other forces, and indeed it is challenging (if not possible) to construct a consistent quantum mechanical theory of particle physics without introducing gravity and the graviton at some point.

For every other force X, we found that there exists a dimensionless coupling constant or strength α_X parameterizing the force, and so we may wish to do the same for gravity. Unfortunately, there isn't one. To see why, let us examine the Newtonian form of the gravitational potential between two masses,

$$U(\vec{r}) \quad = \quad -\frac{G\, M_1\, M_2}{r}, \tag{16}$$

where G is Newton's constant. We know that mass is simply a form of rest energy, and so really the Newtonian potential predicts an attraction between two objects with significant energy[13]. Importantly, note that the potential *grows* in magnitude at larger mass-energy. Therefore, the force of gravity cannot be characterized by a dimensionless number because its magnitude explicitly depends on the energy. This is unlike electromagnetism, the strength of whose interaction is determined by electric charges that are independent of the mass-energy of an object.

The constant giving the characteristic strength of gravitational interactions is in fact G. By dimensional analysis, G has units of (dc^4/energy) where d is distance (this follow from the fact that $M \sim \text{energy}/c^2$). Indeed, we can go one step farther and recall that $d \sim \hbar c/E$ and get that G has dimensions of $\hbar c^5/\text{energy}^2$. We know that the probability of an interaction via the gravitational force scales like some power of G (*i.e.*, it should go to zero if we set Newton's constant to zero). However, probability is dimensionless, we see that G must multiplied by $E^2/(\hbar c^5)$ to obtain something dimensionless. So we find that the probability of an interaction occurring via the gravitational force (via a graviton) should scale like

$$P(E) \quad \propto \quad \left(\frac{GE^2}{\hbar c^5}\right)^n \tag{17}$$

for some positive power n. We see a probability that *grows* with energy. This equation tells us that gravity is not *intrinsically* weak, but because the typical energy of particles in the lab is much smaller

[13] Indeed, Einstein's theory of gravity, which is known as general relativity, shows that massless particles with energy (such as photons) can still experience a classical gravitational force.

than $\sqrt{\hbar c^5/G}$, then the force appears to be very feeble.

Exercise: Check all of the dimensional analysis arguments of this section for yourself.

Hopefully, you can now see an analogy with the weak interactions in Eq. (13) where we also saw that the force got stronger at higher energies. There, we found that the probability grows with energy until the energy of the interaction is sufficiently large to physically produce the force-carrying $W^\pm$ and Z bosons, $E \sim M_W c^2$. At such energies, we find that the strength of the weak force is instead characterized by a dimensionless number α_W. We may expect, then, that at sufficiently high energies $E \sim \sqrt{\hbar c^5/G}$, we have sufficient energy to produce some new ultra-heavy particles that will fix what appears to be a pathological indefinite growth of the probability of interaction with energy.

Here is the problem: within the confines of particle physics, there appears to be no such particle that is consistent with our theory of gravity. In other words, we know that something new should show up in our theory at energies $E \sim \sqrt{\hbar c^5/G}$, but in particle physics there is nothing that can fit the bill and give rise to a sensible theory of gravity at higher energies. There are, however, alternative theories that do remedy these problems. For example, if you allow in addition to fundamental particles the existence of tiny **strings**, then at energies $E \sim \sqrt{\hbar c^5/G}$ you can start producing heavy string states, and these fix the theory. In string theory, we therefore have a consistent quantum mechanical theory of gravity up to arbitrarily high energies, whereas in particle physics the theory is known to break down at some point. This is the famous "incompatibility" of quantum mechanics and gravity.

Fortunately (or unfortunately, depending on how you look at it), the energy at which the gravitational force becomes strong is absolutely enormous. It is known as the **Planck energy**, and it is $E_\mathrm{Pl} \sim 10^{19}$ GeV. It is about a billion Joules, which may not seem huge but we are talking about that much energy confined into a *single particle*! If we directly want to test what happens to gravity at such high energies, we would need a collider that is about a billion billion times more energetic than our current technology. There are suggestions for how we can *indirectly* test new particles with such larger masses, but so far we have seen no positive signals. In any event, since gravity has so little effect on the phenomena in particle physics at the "common" energies we are currently exploring, we will essentially ignore it from now on. This, of course, should not be misinterpreted as my thinking that gravity is uninteresting - in fact, gravity is very interesting, especially if we think about quantum gravity! But, we have a limited amount of time and space to discuss particle physics, so we must leave our discussion of gravity.

It is worth noting that there is one exception to the above argument that could make gravity observable in current and upcoming particle physics experiments. The Newtonian form of the potential and its inverse square law is valid in any space with three spatial dimensions. However, if there existed **extra spatial dimensions**, then some of the gravitational force could "leak" into the extra dimension, making gravity look weaker in our world than it actually is. If these theories are true, it could be that the *true* Planck energy is actually more like 10^4 GeV, which is much closer to the energies of our current experiments. The extra dimension would have to be very tiny; otherwise, we would be able to see that space has four instead of three spatial dimensions! One possible way of thinking about it is that our 3-space-dimensional world is a sort of surface in a higher-dimensional world, and gravity leaks into the extra dimension as shown in Fig. 15.

Many theories such as string theory predict the existence of these extra dimensions, although the extra dimensions in string theory may be too tiny to be detectable in the near (or far) future. While we would need to be very lucky indeed to observe a sign of extra dimensions at a particle physics experiment, it would be an understatement to say that such an observation would fundamentally revolutionize the understanding of our physical world.

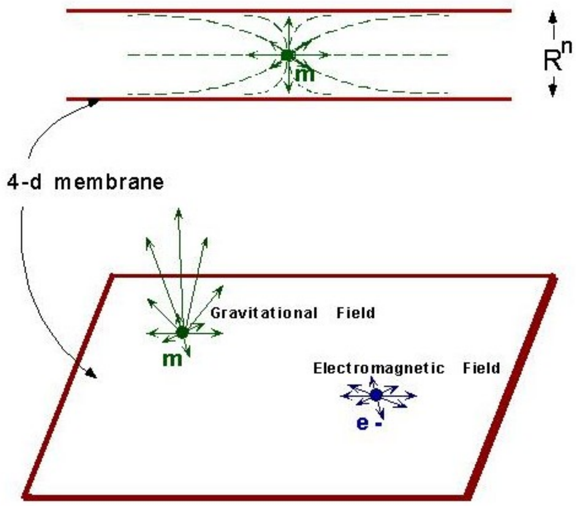

Figure 15: Geometric picture of a world with extra spatial dimensions. Our particles and forces (like electromagnetism) are confined to a 3-spatial-dimensional surface, while the gravitational force can extend into the extra dimensions. This has the effect of making gravity appear weaker than it is (image source: Kapitulnik, Geballe, Beasley Group, Stanford University).

2.5 Summary

Exercise: Create a summary table outlining the properties of the four forces. For each, you should specify the mass of the corresponding mediator particle, the approximate dimensionless coupling constant for each as measured in current experiments, and whether this grows or gets smaller at higher energies.

3 Matter

In the last section, we studied in detail the different fundamental forces in nature. These forces act to attract or repel different kinds of particles. We use "matter" in a broad sense to describe the stuff upon which the forces act (such as electrons and quarks), as opposed to the carriers of the force (such as photons or gluons)[14].

We start with a description of the electron, which is likely to be the most familiar fundamental particle to you. It falls into a broader category known as the **leptons**, which is a fancy word for anything which does not interact via the strong force. We then move on to the matter particles that *do* feel the strong force, which are known as **quarks**. Finally, we must address the thorny fact that particle carrying colour charge are never seen on their own, but rather bound up into colour-neutral objects known as **hadrons**. This category includes familiar particles such as the proton and neutron, as well as more exotic ones such as the pion, and we will discuss the connections between the hadrons and their constituent quark content.

3.1 Charged Leptons

The first subatomic particle to be discovered was the **electron**. It is negatively charged ($Q = -e$) and is electrically attracted to positively charged protons ($Q = +e$), forming neutral atoms. Because the electron is substantially lighter than the proton (with a mass of 511 keV/c^2 compared to the proton's mass of 938 MeV/c^2), it is also much more mobile: the electron can be stripped off of the underlying atom, resulting in the phenomenon of electricity, and it can be moved between atoms and shared within molecules, leading to everything we currently understand about chemistry. The orbital motions of electrons also give rise to magnetic fields. As a result, most of what we know about electromagnetism originates with the dynamics of electrons and their interactions with the electromagnetic field (or, in the language of particle physics, photons).

The reason that the electron in hydrogen is found orbiting the proton at the Bohr radius ($r \sim 10^{-10}$ m) instead of being trapped inside of the nucleus ($r \sim 10^{-15}$ m) is that the electron *does not* feel the strong force. This property of the electron makes it lighter and more agile than nuclei. Because we don't have to worry about the complexities of the strong force, it also makes it easier to calculate properties of the electron in terms of the simpler electromagnetic and weak forces. This makes the electron a desirable testing ground for our various theories of particle physics.

The electron is not alone in its ignorance of the strong force. Indeed, it is one of six particles that are known as **leptons**, the defining characteristic of which is that they do not carry colour charge. Consequently, they do not interact with gluons, which are the mediators of the strong force. Three of the leptons (including the electron) are electrically charged with charge $Q = -e$ (same as the electron). They are in fact identical to one another in all respects but one: the three charged leptons differ in their masses. The charged leptons can be found in the third row of Fig. 7. They are

1. The electron (e), $M_e = 511$ keV/c^2,
2. The muon (μ), $M_\mu = 106$ MeV/c^2,
3. The tau lepton (τ), $M_\tau = 1.78$ GeV/c^2.

We refer to these as the three **flavours** of charged lepton. Their masses also span a large range: the muon is about 200 times heavier than the electron, while the tau is about 20 times heavier than

[14]This distinction is somewhat artificial: for instance, we have already seen that gluons both mediate the strong force *and* also feel the pull of the strong force. However, another distinction between matter particles and force carriers is that matter particles must be produced in pairs, whereas force carriers can be emitted (or absorbed) as a single particle.

the muon. Why there are three flavours and why they have the masses they do remains a mystery. However, the heavier leptons are no longer present in any appreciable quantity in the Universe today. The reason is that the weak force allows heavier flavours to decay into lighter ones (for example, the muon can decay into an electron), and so all of the taus and muons present in the early universe have since decayed into electrons. The same happens for muons and taus produced in high-energy particle colliders. We will examine the interactions of charged leptons, and specifically the weak force that mediates their decays, in Section 5.3.

Finally, we comment that like every other kind of matter particle, charged leptons have corresponding **antiparticles**, which have identical mass and other properties, but an electric charge that is $+e$ instead of $-e$. When particles and antiparticles meet, they can annihilate into energy carried off by photons. The antiparticles of electrons, called **positrons**, were the first antiparticles to be discovered, and this discovery was a triumph of the quantum theory of electromagnetism. By now, we have very precisely tested the properties of antiparticles (such as the mass and magnetic properties of positrons, as well as the properties of "positronium", which is an "atom" made up of an electron and a positron), and all of them have confirmed that positrons have the same properties as electrons (except for the sign of charge) and have served as exquisite tests of the validity of our theory of charged leptons and their interactions. One of the most common types of particle collider experiment is one which collides electrons and positrons to form other types of heavier particles from the collision energy.

3.2 Neutrinos

Neutrinos are funny little particles. They are the only particles that interact *exclusively* via the weak force, and are therefore neutral leptons due to the fact that they are neither electrically charged nor carry colour charge. This means that they hardly interact with anything at all. They are also massless in the Standard Model[15]. Neutrinos are emitted all of the time in decays of various particles via the weak force (for instance, electron antineutrinos are emitted in the beta decay process, $n \to pe^-\bar{\nu}_e$), and are produced in the nuclear fusion process that causes the Sun to burn. Indeed, something like 10 trillion neutrinos from the Sun pass through your body every second. You may in principle worry about the fact that so many particles pass through your body each second: however, a neutrino emitted at solar energies could pass through over a light-year of lead on average without interacting even a single time so there is no reason to be concerned! This gives a sense of just how weak the weak force.

The fact that neutrinos interact so weakly with even dense forms of matter means that its existence is usually inferred from the fact that we *don't* see it after we produce it in an experiment: it passes straight through whatever detectors or apparatus we have without leaving a trace. Indeed, it is the *invisible* nature of the neutrino that first led to its discovery. As mentioned earlier, beta decay (see Fig. 5) produces neutrinos and occurs through the process $n \to pe^-\bar{\nu}_e$. By carefully measuring the energy and momentum of the proton and electron, it became evident that momentum and energy were *not conserved* in the beta decay process! Wolfgang Pauli in 1930 did not want to abandon fundamental principles of physics such as conservation of energy and momentum, and so he postulated the existence of an invisible particle that "carried off" the remaining momentum and energy. This particle soon came to be called the neutrino, and its existence was finally confirmed in 1956 by directly observing a neutrino's interactions with other forms of matter (more specifically, the neutrino's capture on protons in $\bar{\nu}_e p \to ne^+$, leading to neutron and positron emission). Today, neutrinos appear as "invisible" particles at most accelerators and colliders, although special dedicated

[15]It is now known that neutrinos are not exactly massless, but instead have masses of order 0.1 eV/c^2, which apart from being extraordinarily tiny compared to the electron mass already tells us that something is missing from the SM.

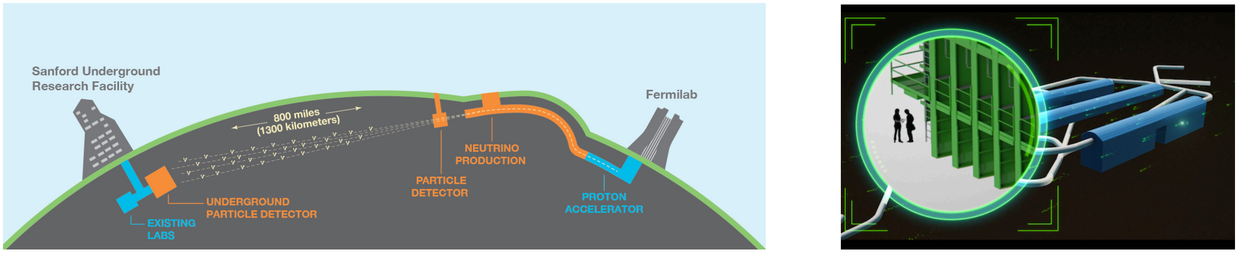

Figure 16: Because neutrinos are so weakly interacting, they pass through the Earth without stopping. The Deep Underground Neutrino Experiment, DUNE (left pane), will take a beam of neutrinos produced at Fermilab (a national lab near Chicago), shoot it through the Earth, and have the neutrinos pass through the Homestake Mine in South Dakota. Such a large number of neutrinos will pass through the Mine that a small number of neutrino scatterings will be observed in a large liquid argon detector (right pane). Image sources: DUNE collaboration.

neutrino experiments have been developed where truly massive amounts of material (thousands of tonnes!) are exposed to intense neutrino beams to look for rare scattering events (see, for example, the planned DUNE experiment in the US in Fig. 16).

As with the charged leptons, there are three flavours of neutrinos: the electron neutrino (ν_e), the muon neutrino (ν_μ), and the tau neutrino (ν_τ). The neutrinos are shown in the bottom row of Fig. 7. Together, each pair of charged and neutral lepton forms a "family" or generation (for example, the muon and muon neutrino form a family), and in the SM each family interacts exclusively with one another (for example, a direct interaction between neutrinos and charged leptons always occurs within the same family: a tau lepton only directly interacts with a tau neutrino and not an electron neutrino, and so on). It is for this reason that beta decay always produces an electron and a corresponding electron (anti)-neutrino. The precise nature of the weak interactions between neutrinos and charged leptons will be elaborated upon in Section 5.3.

An exciting aspect of neutrino physics is that, because they are so hard to detect, they are among the least well-measured particles in the Standard Model. As I have already mentioned, the fact that neutrinos do have a tiny mass already indicates that the SM is incomplete due to its prediction of exactly massless neutrinos, and indeed there may be all kinds of new particles and forces that talk to neutrinos that we have yet to discover. There are reasons, for instance, that there is a connection between neutrinos and the mysterious dark matter that pervades the Universe.

3.3 Quarks

The name "quark" is given to any matter particle that interacts via the strong force; in the language of Section 2.3, we say that quarks carry colour charge. Because of the confining property of the strong interaction, we never see a free quark on its own: instead, they are always bound up into colour-neutral particles called **hadrons**.

The quarks are also electrically charged and therefore interact via the electromagnetic force; they can also decay via the weak force. Quarks are therefore special in that they interact via every type of fundamental force. The strong force typically plays the most important role in the dynamics of quarks, but electromagnetic and weak interactions can also be important: for example, the beta decay of a nucleus is due to a weak interaction process involving the constituent quarks, while the electromagnetic interaction of an electron with the quarks inside of the proton gives rise to the formation of atoms.

There are two classes of quarks, which we call **up-type** and **down-type**. The up-type quarks have electric charge $+2e/3$, while the down-type quarks have electric charge $-e/3$. The antiquarks are the same as the quarks but with opposite charges ($-2e/3$ and $+e/3$, respectively). Notice that $Q_d - Q_u = -e$, while $Q_e - Q_\nu = -e$; in other words, the charge splitting between up and down-type quarks is the same as the charge splitting between charged leptons and neutrinos. Since we paired each charged lepton with a neutrino to form a "family" (for instance, the tau lepton and the tau neutrino formed a family pair), we can also make pairs out of the up- and down-type quarks. Just as with the leptons, there are three families of quarks: since there is an up- and down-type quark in each family, this gives six quark **flavours** in total. The families are:

1. First family: up quark ($M_u = 2.3$ MeV$/c^2$, $Q_u = +2e/3$) and down quark ($M_d = 4.8$ MeV$/c^2$, $Q_d = -e/3$).
2. Second family: charm quark ($M_c = 1.29$ GeV$/c^2$, $Q_c = +2e/3$) and strange quark ($M_s = 95$ MeV$/c^2$, $Q_s = -e/3$).
3. Third family: top quark ($M_t = 173$ GeV$/c^2$, $Q_t = +2e/3$) and bottom quark ($M_b = 4.18$ GeV$/c^2$, $Q_b = -e/3$). These quarks are also sometimes referred to instead as truth and, especially in Europe, beauty.

The quarks are found in the top two rows of Fig. 7. As you can see, the quarks have rather whimsical names[16]. The quark masses span many orders of magnitude, from the up and down quarks with masses close to the electron mass all the way to the top quark, which has is a single elementary particle with a mass comparable to a gold atom.

The heavier quarks can undergo weak-force-mediated decays into the lighter-flavour quarks, and so nowadays nearly everything in the Universe is made up exclusively of up and down quarks. However, heavy-flavour quarks are created in high-energy particle collisions and were present in the early universe. There's also something special about the fact that there are three families of quark: a theory with three families of quark is allowed to have a property whereby matter and antimatter have slightly different properties. This could, therefore, have something to do with the fact that everything we see in the universe is made of particles and not antiparticles. However, we have no fundamental reason for why there are three families of quark, nor why the masses for the quarks are so widely separated. As far as we can tell, these things are just true, but it would be interesting to have some theoretical explanation for why the SM particles have this organization.

As I have stated repeatedly, we never see quarks (or gluons) on their own because of the confining property of the strong force. In Sec. 3.4, we will take a look at the composite bound states of quarks that do appear in nature, including our old friends the proton and neutron. However, we do see *evidence* of quarks and gluons: they can exist as free particles very briefly in very high energy particle collisions (remember that the strength of the strong force decreases with higher energies, and so quarks can exist temporarily as free particles during the collision), and the imprint of those momentarily free quarks and gluons are seen as sprays of colour-neutral pointing in the same direction as the original quark or gluon. We examine this more in the notes on collider physics.

Exercise: We observe protons to have charge $+e$ and neutrons to have charge 0, and each is made exclusively of quarks (rather than antiquarks). What is the minimum number of quarks that could possibly make up a proton and a neutron? Which combinations of up- and down-type quarks could make up a proton and neutron? Of the six flavours of quarks, which do you think make up the proton and neutron? Why?

[16] Many of the quark names have their origins in a time before we knew of the existence of quarks: for example, a certain class of hadrons exhibited strange behaviour and so were (creatively) called "strange particles". We now know that the characteristic of a strange particle is that it is made up of at least one strange quark.

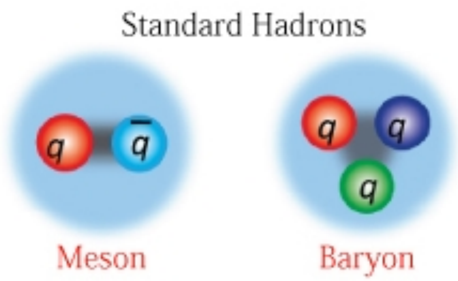

Figure 17: Illustration of the two different types of colour-free hadrons: mesons (made up of a quark of a particular colour and an antiquark with the corresponding anti-colour, in this case red and anti-red) and baryons (made up of three quarks, one of each colour of green, red, and blue). Image source: Quantum Diaries.

3.4 Hadrons

Hadrons are particles that carry no colour charge, but are made up of constituent particles (namely quarks) that themselves carry colour charge. The force holding together the hadrons is the strong force, and in a very naïve way we can think of hadrons as bound states in the same way that atoms are bound states. The main difference is that it is relatively easy to kick an electron out of orbit around an atom (or even to knock a proton or neutron out of the nucleus), while it is impossible to completely separate out the constituent quarks: the energy needed to pull a quark out of a hadron is larger than the mass-energy needed to produce an additional quark-antiquark pair, so additional quarks and antiquarks are created as needed to ensure that we are only left with colour-neutral hadrons at the end of the day.

Every quark except for the top quark can be found inside of hadrons. The reason why top quarks are *not* found in hadrons is that they decay so quickly into lighter flavour quarks that the decay occurs faster than the characteristic time scale for hadron formation, $\tau \approx 10^{-23}$ s.

Exercise: We saw in Sec. 2.3 that the strong interactions become strong at energies $E \sim 0.2$ GeV. Using Planck's constant, we can define a corresponding time scale $t = \hbar/E$ of hadron formation. Numerically compute this value, compare this to the top quark decay lifetime and explain why top quarks do not form into hadrons.

As described in the section on the strong force, there are two ways we can get colour-neutral particles out of combining various quarks carrying colour charge. Each of these corresponds to a separate class of hadrons (see Fig. 17):

1. **Mesons:** These are particles made up of a quark and an anti-quark. For example, the colour of a red quark can be cancelled by an anti-red anti-quark, and so this forms a colour-free bound state.

2. **Baryons:** These are particles made up of three quarks, each of which carries a different colour. For example, a red quark + a blue quark + a green quark will form a colour-neutral combination. This is due to the weirdness of the strong interactions and the behaviour of colour charge. Similarly, an antibaryon is made of three antiquarks, each of which carries a different anti-colour that combines to once again form a colour-neutral bound state.

Because mesons are made of a quark-antiquark pair, the quark-antiquark pair can annihilate into lighter particles. From the point of view of the meson, it appears as though it is disintegrating (because the constituents literally blow themselves apart) As a result, all of the mesons are unstable. By contrast, the lightest baryon (the proton) is absolutely stable in the SM[17], and its partner the neutron is also very long-lived inside of nuclei. Therefore we are most familiar with baryons and we will turn to these first.

3.4.1 Baryons

The most commonly known baryons are the proton and the neutron. You may have already worked out the quark composition of a proton and neutron above. The proton and neutron are made of three quarks each and have electric charges $Q_p = +e$ and $Q_n = 0$. We can therefore see that a proton is made of two up-type quarks and one down-type quark, whereas a neutron is made of two down-type quarks and one up-type quark. Indeed, as the lightest of all baryons, the proton and neutron are made entirely of the lightest flavours of quark, and hence we say that protons are made of uud while neutrons are made of udd. The antibaryons are made of the corresponding antiquarks: the antiproton $\bar{p}$ is made of $\bar{u}\bar{u}\bar{d}$, etc.

We can also make baryons with different flavours of quarks. For example, the strange baryons have at least one strange quark in them: an example is the Λ^0 which is made of uds. Similarly, there exist baryons with charm and bottom quarks.

Let us examine the properties of the familiar baryons. The masses of the proton and neutron and nearly identical, with $M_p = 938$ MeV$/c^2$, and $M_n - M_p \approx 1$ MeV$/c^2 \sim 10^{-3} M_p$. We can immediately see two things. First, the proton and neutron masses are *not* simply the sum of the masses of the constituent quarks. Indeed, the sum of quark masses is ~ 10 MeV$/c^2$, meaning that roughly 99% of the mass of the proton and neutron have *nothing* to do with the underlying quark masses! So where does the mass come from?

We can think of hadron is a bunch of quarks held together by a **colour field** mediated by gluon exchange, and there exists kinetic energy as the quarks inside of the hadron orbit one another. It so happens that there is *a lot* of energy stored in the colour field inside of a hadron (and a comparable amount in the quark kinetic energy), and this energy manifests itself as a mass for the baryon according to $E = Mc^2$ in the baryon's rest frame. We find that the strong force itself is responsible for the vast majority of the mass of the proton and the neutron.

The second interesting thing we see is that the proton and neutron are very close in mass. The reason is related to what we found in the immediate paragraph: the u and d quark masses play almost a negligible role in determining the p and n masses. Indeed, if the u and d quarks were massless, the protons and neutrons would have exactly the same mass[18].

[17] We are fortunate that the proton is stable: otherwise, atoms (and, by extension, each of us) would not exist! Extensions of the Standard Model feature new forces that can allow protons to decay, and so the fact that we have not yet observed protons to decay allows us to constrain the existence of these new forces. For example, the decay lifetime of the proton should not be much shorter than the age of the Universe ($\sim$ 14 billion years), or else there would not be many protons left today. In fact, the absence of observed proton decay in large quantities of matter studied in experiments puts a much more stringent lower limit of 4×10^{29} years on the proton lifetime.

[18] This is a manifestation of a symmetry known as isospin symmetry, which is roughly speaking a symmetry under interchange of protons and neutrons (or, equivalently, up and down quarks). This turns out to be an extremely important symmetry for understanding the organizations of hadrons and their behaviour under the weak interactions, and helped physicists first elucidate the quark structure of hadrons. Now that we already know there exist quarks, though, the isospin symmetry and other patterns of masses and properties among hadrons are understood as simple consequences of how quarks are distributed among the various hadrons.

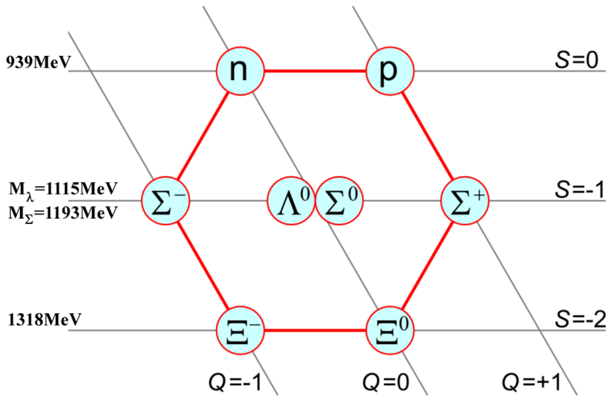

Figure 18: Lowest-energy baryon states in the Standard Model (with the proton and neutron in the top row). The baryons are made up exclusively of up, down, and strange quarks. The horizontal lines are distinguished by the number of strange quarks (S) inside of the baryon, while the diagonal lines indicate the electric charge, Q, of the baryon. This is known as the "baryon octet". There exist many, many other baryons with different combinations of constituent quarks and their angular momenta. The masses are shown in units of MeV$/c^2$ (image source: Wikipedia).

Strange and charm baryons are still heavier than the proton and neutron, in part due to the fact that the strange and charm quarks are substantially heavier than the up and down. For example, the $\Lambda^0(uds)$ has a mass of 1.115 GeV$/c^2$, which is about 75 MeV$/c^2$ heavier than the proton and neutron. This is largely accounted for by the heavier s quark mass. Similarly, the charm baryon $\Lambda_c^+(ucd)$ has a mass 2.286 GeV$/c^2$, which is heavier than the proton and neutron by approximately the charm quark mass. The heavier strange and charm baryons can decay to the lighter proton and neutron via the weak interaction, with the charm or strange quark changing into an up or down quark. The proton is absolutely stable, while the neutron can decay through the beta decay process $n \to p^+ e^- \bar{\nu}_e$, although its lifetime inside of a nucleus is very long such that we often see neutrons inside of stable nuclei.

There are a plethora of baryons. One way to obtain different baryons is to combine different flavours of quarks (see Fig. 18 for an example). It is also possible for the bound states of the same quarks to be in different energy states (think of the hydrogen atom and its different electron orbitals) depending on the alignment of the quarks and their relative momentum. Originally, this "particle zoo" of hadrons was a mystery to particle physics researchers, who could not understand why there existed so many seemingly fundamental particles. We now know that this zoo is purely due to the myriad ways quarks can be combined into hadrons.

Exercise: What is the quark content of the Ξ^0 baryon in Fig. 18? Why is it heavier than the Λ^0 baryon?

3.4.2 Mesons

The other type of meson is a quark-antiquark bound state. This is colour neutral because the colours of the quark and antiquark state can cancel out (for example, a blue quark can be paired with an anti-blue anti-quark in the meson). There are three possible electric charges for mesons:

- $Q = +e$, bound state of an up-type and anti-down-type quark such as $u\bar{d}$ (or other flavours, such as $c\bar{s}$);

- $Q = 0$, bound state of either up-type and anti-up-type quarks or down-type and anti-down-type quark (such as $d\bar{s}$);
- $Q = -e$, bound state of down-type and anti-up-type quark such as $d\bar{u}$ (or other flavours, such as $s\bar{c}$);

As mentioned earlier, because mesons are bound states of a quark and an antiquark, the quark and antiquark can annihilate leading to meson decay. Mesons can decay to either lower-mass hadrons, or else to leptons. For example, the ϕ^0 meson (which is a bound state of $s\bar{s}$) can decay to $K^+(u\bar{s}) + K^-(\bar{u}s)$ mesons.

Exercise: Explain in your own words why the annihilation of a quark-antiquark pair can appear as a decay to a distant observer if the quark-antiquark pair are in a bound-state meson. How far away from the meson do you have to be for it to look like a decay instead of a quark-antiquark pair annihilation? If it helps, you can think of the meson bound state like a proton-electron bound state inside of an atom.

The most important set of mesons are the lightest mesons, which are known as **pions**. The pions come in charged varieties, $\pi^+(u\bar{d})$, $\pi^-(\bar{u}d)$, as well as a neutral variety (π^0, which is a combination of $u\bar{u}$ and $d\bar{d}$). The π^- is the antiparticle of the π^+, while the π^0 is its own antiparticle (note that there is no good distinction of "matter" vs. "antimatter" for mesons because each meson is made of both a quark and an antiquark). Due to the fact that up and down quarks are nearly massless, the pions have almost the same mass ($M_{\pi^+} = 140$ MeV/c^2, $M_{\pi^0} = 135$ MeV/c^2)[19]. Most important is the fact that they are *much lighter* than the lightest baryons, which have masses $\sim$ GeV. It turns out that, if the up and down quarks were exactly massless, the pions would also be exactly massless. As the lightest mesons, pions are often the end products in the decays of heavier hadrons.

Since the pions are the lightest hadrons, they *cannot* decay into other hadrons. Therefore, they must decay into other, non-strongly-interacting particles. For example, the π^+ decays predominantly into an anti-muon and a muon neutrino (via the weak force), while the π^0 decays predominantly into two photons (via the electromagnetic force).

It turns out that pions also play an additional role as the force carriers of the **residual strong force** (illustrated in Fig. 14). Hadrons do not carry colour charge themselves and the colour force is confined *within* the hadrons, which means that two hadrons cannot interact directly by exchanging gluons. However, it is possible for the quarks inside of one hadron to emit a quark-antiquark pair that forms into a pion, which as a colour-free object *can* exist independently outside of the hadron. Thus, hadrons can interact via the exchange of pions, and the pions act as a force carrier much in the same way that photons and gluons do (this is true even though the pions are not fundamental particles but in turn made of quarks). This residual force, also known as the **Yukawa force**, is responsible for the attractive force holding protons and neutrons together inside of the nucleus: were it not for pions and the residual strong force, the strong positive charge concentrated inside of nuclei would result in the nucleus blowing itself apart[20].

[19] The mass difference between the charged and neutral pions is largely due to the fact that the charged pion interacts via the electromagnetic force while the neutral pion does not.

[20] As an interesting historical aside, Yukawa predicted in 1935 the existence of a force-carrying "meson" particle to explain the binding of protons inside of nuclei. In 1936, a new particle was discovered called the "mu meson", and it was assumed to be the particle that Yukawa predicted. Unfortunately, it seemed to exhibit none of the *properties* expected of Yukawa's particle (such as interacting strongly with nuclei); the "true" Yukawa particles, the "pi mesons", were not discovered until 1947. It was then realized that the mu meson was indeed a lepton (now known as the muon) and had nothing to do with the residual strong force, while the pi meson we now call a pion. The fact that no one expected a heavier copy of the electron to exist led to the famous exclamation of Raby regarding the muon: "Who ordered that?!"

Just like the baryons, mesons come in a whole host of different types depending on the flavours of the constituent quarks and the arrangements of the quarks inside of the meson (such as the angular momentum, etc.). For example, the strange mesons contain one strange quark and are known as **kaons**, while the charm mesons contain one charm quark and are known as ***D*-mesons**. Even more so than for the baryons, the properties and structures of the mesons can be derived from the underlying quark theory. As this lies beyond the scope of our study, however, we will content ourselves to this cursory look at the mesons: you will return to this subject in more detail if your research requires a more detailed knowledge of mesons and hadronic physics.

Exericise: Which do you expect to be heavier: kaons or charm mesons? Why?

3.5 Summary

Having now met all of the different types of elementary particles, we see that they tend to come in three families. Indeed, we can associate the electron, electron neutrino, up, and down quarks as a "first family" (or first generation); the muon, muon neutrino, strange, and charm quarks as a "second family" (or second generation); and, finally, the tau, tau neutrino, top and bottom quarks as a "third family" (or third generation). Indeed, it was discovered in the 1970s that the mathematical self consistency of the physical description of the Standard Model demands that particles come in complete families made up of a charged lepton, a neutrino, an up-type quark, and a down-type quark; however, we do not understand why there are *three* such families.

4 Particle Interactions & Feynman Diagrams

So far, we have learned about all of the different kinds of particles and forces in the SM, and I have given a largely qualitative overview of the various kinds of particle interactions. In this section, we will look at particle interactions in more detail by studying precisely how different forces cause particles to influence each others' dynamics. To understand the full story, you will need to study the subject of **quantum field theory (QFT)**, which is the foundation on which our modern theory of particle physics rests. The basic idea of QFT can be understood from our picture of electromagnetism: the electromagnetic force is manifest in electric and magnetic fields, and yet these fields themselves are made up of tiny lumps of energy called "photons". Photons are the most basic component of the field, and the classical behaviour of the electromagnetic field (for example, the light we see from a light bulb) is simply the result of the collective behaviour of an enormous number of photons. Every other kind of particle has an associated field too: there is an "electron field" that fills space, and the excitations of this field are lumps of matter called electrons. This approach allows us to put matter and forces on the same footing and calculate the properties of interactions between different kinds of particles.

Quantum field theory is a beautiful and incredibly powerful tool for studying various phenomena in physics, but it is advanced enough that we will not be able to derive the main results of particle physics using it; you will need to wait until you have studied quantum mechanics, classical mechanics, and relativity in great quantitative detail before you can embark on learning it in any meaningful way. Fortunately, many of the main results of QFT can be understood in terms of elementary particles undergoing pointlike interactions (meaning that particles interact by colliding at a particular point). We can therefore apply some of our classical reasoning to simple problems in QFT. This approach is due to Richard Feynman, who invented a new way of thinking about quantum mechanics and, with it, a pictorial method for describing particle interactions. These pictures, known as **Feynman diagrams**, allow us to visualize the way that particle interactions occur, to specify what kinds of interactions are allowed/forbidden, and thereby to understand the basics of particle physics. The pictures also encode the ingredients used to perform calculations of particle interaction probabilities and rates, although we will not delve into the quantitative aspects here. Feynman diagrams are used in every aspect of particle physics, and learning the diagrammatic forms of each of the forces will be crucial to your understanding of the dynamics and scattering of elementary particles.

We start this section with a detour into regular quantum mechanics, specifically a discussion of the **double-slit experiment**. The reason is that the results of the double-slit experiment motivate a particular way of thinking about physical processes in quantum mechanics that leads directly to the use of Feynman diagrams to model the effects of particle interactions. We will then use a simple example, namely a theory with an electron, a positron, and a photon, to illustrate how Feynman diagrams can be used to represent scattering processes. In Section 5, we will turn to the particle interactions due to the other forces in the Standard Model.

4.1 Quantum Mechanics and the Double-Slit Experiment

It is often stated that, in quantum mechanics, objects have both "particle-like" and "wave-like" behaviour (sometimes referred to as **wave-particle duality**). Perhaps you have heard this in an early course or unit on modern physics, or in popular science articles. What does this statement mean?

4.1.1 Wave-Particle Duality

Prior to the development of quantum mechanics, people thought that all of matter was divided into two types of objects: particles and waves. Particles were like billiard balls or pellets that could be shot out of a tube: the particle has a definite position and momentum at any given time, and it follows a trajectory that you can calculate using the principles of Newtonian mechanics ($\vec{F} = M\vec{a}$). Atoms, their constituents, and everything they were made of were thought to be "particles". In contrast, "waves" were understood as the collective motion of a system over an extended area: think about the ripples travelling across a pond, or the electromagnetic fields that give rise to light rays. Because waves are extended over a region of space, they can undergo more complex motions such as diffraction and interference. It was traditionally believed that something was *either* a particle or wave.

One of the principal observations driving the development of quantum mechanics was the finding that light (which physicists originally thought were waves) actually displayed particle-like behaviour. The most convincing indication of the particle-like behaviour of light is in the photoelectric effect. We illustrated this effect in Fig. 8, but briefly review it here. The photoelectric effect is the finding that, when light is impinged on a metal, the ejection of electrons at the surface of the metal and the kinetic energy of the ejected electrons is determined by the *colour*[21] (wavelength, or equivalently, frequency) of the light, rather than the total amount of energy in light hitting the surface. The reason is that light is made up of irreducible lumps (now called photons) whose energy is determined by the frequency of the light ($E = h\nu$, where h is Planck's constant and ν is the frequency of the light; see Appendix A.1). Since each electron interacts with at most one photon at a time, the electron can only be ejected if the photon carries enough energy to kick it out of the atom. By contrast, the intensity of light (*i.e.,* how bright the light is) is simply a matter of how many photons there are in the light ray. The light could be very bright, but if each photon is insufficient to eject an electron from an atom, there will be no current associated with the ejected electrons no matter how intense the light. The photoelectric effect shows that wave-like phenomena such as light can also exhibit particle-like behaviour, such as energy being localized to a small region of space in lumps and the lumps able to undergo point-like interaction with other forms of particles, such as electrons.

If waves could exhibit particle-like behaviour, then can particles exhibit wave-like behaviour? In particular, if particles such as electrons posess a corresponding wavelength (or, equivalently, are spatially extended over some area), then perhaps they can also undergo wave phenomena such as diffraction, which is the bending of light around the edges of objects? This was first hypothesized by de Broglie and confirmed in the Davisson-Germer experiment which demonstrated that electrons do, in fact, undergo diffraction. The reason why we don't normally see wave-like behaviour in electrons is because their wavelengths are incredibly tiny compared to light and so our eyes are not sensitive to this "fuzziness"; however, experiments that are specially designed to look for such short-wavelength behaviour can indeed see that electrons (and other types of particles) behave as waves. As quantum mechanics developed, we see other wave-like behaviour for electrons: for example, the fact that electrons are distributed over **orbitals** inside of atoms tells us that they are not, quite, pointlike particles.

4.1.2 The Double-Slit Experiment

If we accept that matter can behave both like waves and particles, an important question is how do we reconcile these two pictures? Is matter localized in one place, or extended over a non-zero area? One of the clearest ways of untangling this knot is through the famous **double-slit experiment**

[21] Here, we are back to talking about actual colour of visible light, not the charge associated with the strong force.

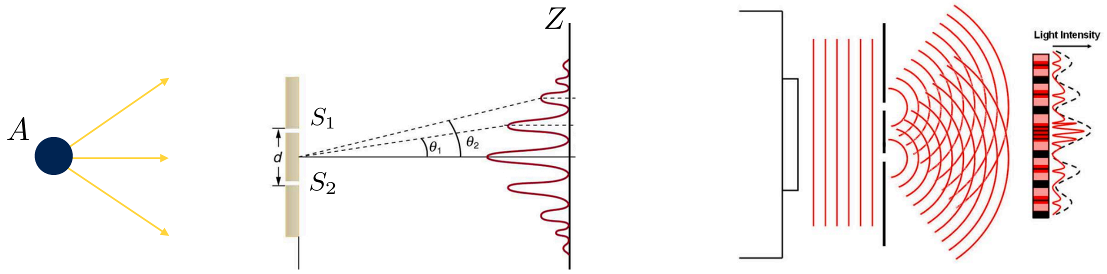

Figure 19: The double-slit experiment. Light is shone from a source A on an opaque screen with slits S_1 and S_2, and then reaches a screen Z. The light arrives as a right-going wave at the slits, and the interference of the parts of the wave passing through slits S_1 and S_2 gives rise to a diffraction pattern at Z (left pane, image source: Rice Open CNX). The diffraction can be thought of as due to the fact that S_1 and S_2 act as their own point sources for new circular waves, with the peaks and troughs of the diffracted waves adding constructively or destructively to give the resulting diffraction pattern (right pane, image source: CK12).

(see Fig. 19). In this setup, light is produced at a source A and shone through an opaque material that has two slits (S_1 and S_2) cut in it. A screen (Z) is used to observe the light on the other side of the slits. When light is shone through the double slits, an interference pattern appears on the screen due to the diffraction of light. This appears to confirm the wave picture of light!

Now, consider that the source A is reduced in intensity such that the minimum possible amount of light is produced. In other words, only one photon at a time is emitted from the source A. In this case, we detect at each moment an individual "hit" of light located at a single point on the screen. This appears to confirm the particle picture of light! However, if we wait long enough, and record the sum total of all of the (many, many) photon hits, we recover the original interference pattern (illustrated in Fig. 20). This tells us two things:

1. Light is produced in irreducible clumps called photons. This is why we only see one hit at a time on the screen when the intensity of the light is reduced.

2. The interference pattern is still present when we look at only a single photon traversing the slits at a time. Even though the photons are going one at a time, they "know" about the wave pattern of the light propagation. Therefore, the wave property is present even at the level of individual particles, and a single particle "knows" about the diffraction through both slits.

How can we reconcile the particle and wave properties? In quantum mechanics, we say that each *individual* photon possesses a wave-like property called a **wavefunction**, typically denoted $\psi(x)$ (where x is the coordinate along the screen and ψ is a different complex number at each point x). Since $\psi(x)$ is non-zero at multiple values of x, it is extended in space and can undergo wave-like behaviour, leading to diffraction through the double slits. However, we know that the individual photon only hits the screen at one place! To map from this extended wave-like object $\psi(s)$ to the fact that we only observe a photon at one particular point, we interpret $\psi(x)$ as being related to the *probability* of finding the photon at point x. In fact, the probability of the photon being observed at the point x is given by $\mathcal{P}(x) = |\psi(x)|^2$. The "wave" in wave-particle duality is essentially a probability wave telling us where we are likely to find the particle if we measure it. This reconciles the fact that each photon is only observed at one place, while the distribution of photons obeys the interference relation predicted by the diffraction of a wave (in this case, the wave is $\psi(x)$).

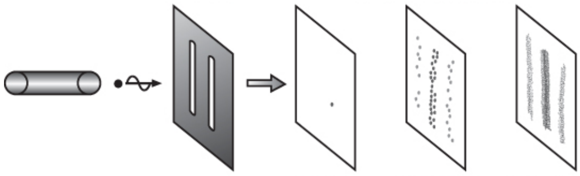

Figure 20: The double-slit experiment, continued. If we dim the light source A so that it only emits one photon at a time (and only one photon hits Z at a time at a particular location), the pattern resulting from the passage of many photons at different times still forms the double-slit diffraction pattern at Z. This tells us that each photon "knows" about the wave property of light, and in particular is sensitive to passing through both slits even though it only hits the screen at one point (image source: Perimeter Institute).

Let us see how this works in practice by returning to the double-slit problem. We can see that, if a photon is to hit the screen at point x, it can pass through either slit S_1 or S_2. The path through S_1 gives a contribution to the wavefunction at point x which we call $\psi_1(x)$; remember, this is simply a number[22]. Similarly, the path through S_2 gives a contribution to the wavefunction at point x which we call $\psi_2(x)$. If we see a flash at point x, then we know that the photon *could* have come from either S_1 or S_2, so we *add* the wavefunctions: $\psi(x) = \psi_1(x) + \psi_2(x)$. The probability is then found by $\mathcal{P}(x) = |\psi_1(x) + \psi_2(x)|^2$. If the path from S_1 gives rise to a positive number evaluated at x while the path of the particle-wave from S_2 gives rise to a negative number at x, the two can cancel to give a suppressed probability of detecting the photon at x (this is known as **destructive interference**). At some point, the two exactly cancel which is why we see points of zero light intensity along the screen. Conversely, if the particle-wave passing through each of S_1 and S_2 are both positive (or both negative), they add and give rise to an even bigger probability of finding the photon at x (this is known as **constructive interference**). This behaviour based on interference of the contributions of the two paths is shown in Fig. 21, with a blue colour representing a negative value of $\psi(x)$ and red representing a positive value of $\psi(x)$ associated with a given path. The above description explains the characteristic peak and trough structure of the interference pattern in Fig. 19: the reason is that the photon in some sense goes through *both* slits (or, rather, has contributions from going through both slits), and these two contributions can interfere to give either an enhanced probability of detecting the photon or a suppressed probability for detecting the photon at a given point x.

To recap, we have seen that the wave-like nature of particles means that, even though each particle only hits the screen one at a time, it can proceed through multiple paths. Each path I is assigned a wavefunction $\psi_I(x)$, and these are summed together and then squared when computing the total probability. Because the wavefunction can be positive or negative, this gives rise to an interference pattern *even though we are only looking at one particle at a time!*. The key point is that the wave-like behaviour stems from the several possible paths that the photon takes: since we don't know which one it took, it in some sense gets contributions from all of them.

Exercise: What do you think happens if we close S_1 and leave only S_2 open? Does the interference pattern stay, or go away?

[22] In fact, it is a complex number but that is not particularly relevant for our discussion.

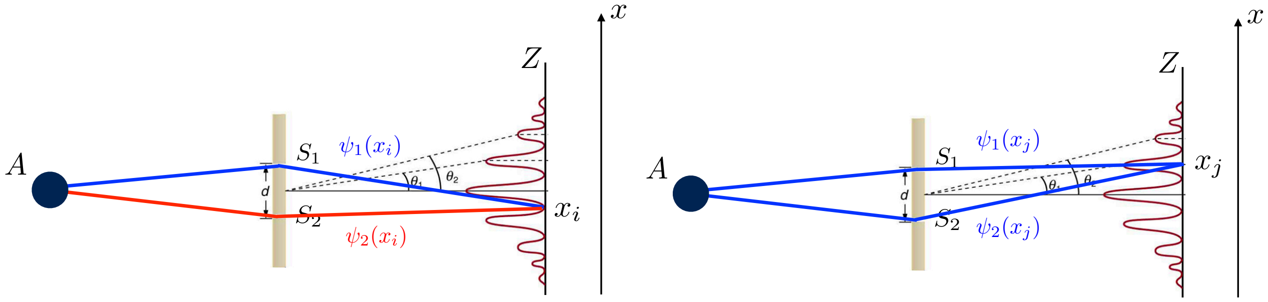

Figure 21: The particle-wave interpretation of the double-slit experiment. Each photon is described by a **wavefunction**, $\psi(x)$, which is associated with the probability of detecting the photon at point x, $\mathcal{P}(x) = |\psi(x)|^2$. Since the photon can travel through either slit S_1 or S_2, there is a contribution to the wavefunction from each path, $\psi(x) = \psi_1(x) + \psi_2(x)$. In the left pane, we consider a photon detected hitting the screen at point x_i. The wavefunction for the path going through S_1 to point x_i has opposite sign from that going through S_2 to x_i, resulting in *destructive* interference and a cancellation in the probability $|\psi_1(x_i)+\psi_2(x_i)|^2 \ll |\psi_1(x_i)|^2$, $|\psi_2(x_i)|^2$. In the right pane, we consider a photon detected hitting the screen at point x_j. Here, both the wavefunction from paths S_1 and S_2 have the same sign, and so there is an enhanced probability of seeing the photon there due to *constructive* interference in $|\psi_1(x_j) + \psi_2(x_j)|^2$.

4.1.3 Tagging Paths in the Double-Slit Experiment

The odd behaviour of wave-particle duality stems from the fact that there are many paths the photon can take. What happens if instead we *know* which path the photon takes? Would the interference pattern change? One way we can test this is to add an extra component to our experiment: two small detectors, one of which is located by each slit. Let us call them D_1 and D_2 associated with slits S_1 and S_2, respectively. Each detector gives off a flash when a photon passes through, but the photon's trajectory is otherwise unchanged. Therefore, we can tell whether the photon goes through slit S_1 by seeing whether a flash goes off in D_1 at the same time. Similarly, we can tell if the photon goes through S_2 if we see a flash in D_2. The set-up is illustrated in Fig. 22.

If we re-run the experiment in this new configuration, we see a dramatic change in the pattern of photons hitting the screen. In particular, the interference pattern changes and the distribution of photon hits on the screen Z resembles two separate single slits. The reason is this: the interference pattern was a result of the fact that the photon took "both paths", in some sense (or rather, both paths contribute to the final probability of where it lands on the screen). Since we added the detectors D_1 and D_2, though, we know on a photon-by-photon basis exactly which path a particular photon took. For example, in Fig. 23 I illustrate the case where a photon went through slit S_1 (firing D_1 in the process) and lands at point x. Therefore, *only* the path through S_1 contributes to determining where that photon lands, and the wavefunction is $\psi(x) = \psi_1(x)$. $\psi_2(x)$ does *not* contribute because we observed that the photon did not pass through the slit S_2.

This seems a bit like magic. How did the photon "know" that we had put detectors there, and which detector fired? Classically, this is a mystery because we can construct our detector in such a way as to minimally disturb the photon and so there is no way that the photon know that we have observed it. In quantum mechanics, however, we find that the *only* to observe the photon is by disturbing its state. By measuring it with detector D_1, we have now changed its state so that it *only* goes through slit S_1 with no contribution from the S_2 path, and consequently the interference pattern disappears. How this happens precisely on a microscopic level (and the interpretation thereof) is

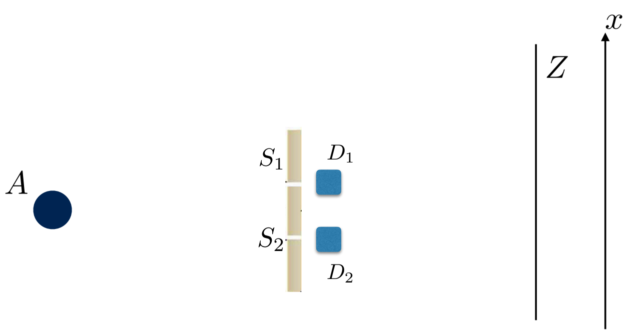

Figure 22: The double-slit experiment with an additional photon detector in front of each slit. The detector D_i fires if the photon passes through slit S_i.

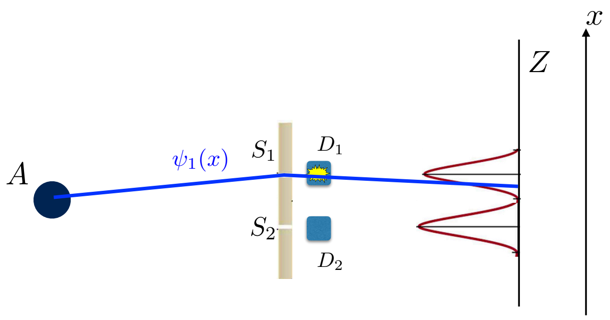

Figure 23: The double-slit experiment with an additional photon detector in front of each slit. The detector D_i fires if the photon passes through slit S_i. We then see a change in the interference pattern: for example, if D_1 fires then we *know* that the photon passed through slit S_1. In this case, there is **no contribution** from the path through S_2 because we know the photon did not go that way. As a result, we have $\psi(x) = \psi_1(x)$ and consequently there is no interference in the pattern of hits on the screen. Instead, it looks like a simple sum over two single slit distributions (one arising from photons passing through S_1 and firing D_1, and the other coming from photons passing through S_2 and firing D_2.

subject to much debate, with ideas like "wavefunction collapse", "decoherence", and "entanglement" relevant to the discussion. However, everyone agrees with how we *use* quantum mechanics, namely that if we observe the photon to pass through detector D_1, then only the part of the wavefunction passing through slit S_1 (namely $\psi_1(x)$) contributes to the probability of measuring the photon at point x. If, however, we can't distinguish from our measurement whether the photon passed through S_1 or S_2 like in our original set-up, we must add the contributions from both S_1 and S_2, giving rise to interference effects[23]. In summary, the trajectory a particle takes (and, accordingly, its position) is not known with certainty until we measure it: unless we have detectors that tells which slit the photon went through, it in some sense went through both.

[23]This is in, some sense, an operational definition of quantum mechanics. We are not specifying the "how" or "why" the quantum state changes through measurement; however, we can specify what the state looks like after observing a certain experimental outcome. This is sufficient for solving particle physics and most other types of physics problems, but the question of what is going on at a "deeper level" in quantum mechanical measurements and systems is the subject of research in the area of quantum foundations.

4.1.4 Quantum Mechanics, Observations, and Interference

We can now summarize a general rule in quantum mechanics: **if we start with a system in state A and wish to compute the probability that it ends up in state B, we must sum over all possible ways for the system to proceed from state $A \to B$ in an indistinguishable manner**[24]. In other words, if we can't tell which "path" the system took from A to B, we have to add together/interfere all such paths in computing the total probability of going from $A \to B$. This is the major point where quantum mechanics is different from classical mechanics: classical mechanics says that a system always evolves along a *particular* path from A to B, whereas in quantum mechanics you must simultaneously consider the effects of many such paths. If, instead, we *observe* that the system follows a *particular* path (or subset of paths), then only those paths contribute to the probability of going from state $A \to B$.

Before we move on to the next section, you may be wondering what all of this has to do with particle physics, which appears at face value quite different from photons traversing slits in a screen. However, precisely the same principles of quantum mechanics apply to particle interactions, even if they do so in a somewhat more abstract way. Suppose we are considering a collision at the Large Hadron Collider: we start off with two colliding protons, and let's say we want to compute the probability that the outcome of this collision is the production of an electron-positron pair. There are many different particle interactions that could lead to this outcome (for example, this could proceed through either the electromagnetic or weak interactions). In this case state A is the proton-proton initial state, state B is the final e^+e^- pair, and the "paths" are the different forces that could possible mediate the interactions turning a pair of protons into an e^+e^- pair. What quantum mechanics tells us is that we have to add together and interfere the quantitative contributions of each of these interactions when computing the probability for it to occur. In order to do this, we need a simple way of systematically figuring out which forces can be responsible for taking our initial state and turning it into our final state. This is the formalism of **Feynman diagrams**, to which we now turn.

Exercise: Suppose we have a three-slit experiment. A photon produced at source A can go through one of three slits: S_1, S_2, or S_3. We also have two detectors, D_A and D_B. D_A gives off a flash of light if the photon passes through S_1, while D_B gives off a flash of light if the photon passes through EITHER S_2 or S_3. Is there an interference pattern formed from the events where D_A gives off a flash of light? What about where D_B gives off a flash of light?

4.2 Quantum Mechanics of Particle Interactions & Feynman Diagrams

4.2.1 Definition of Particle-Physics Observables

We now wish to extend the findings of the last section to our study of particle-physics interactions. Because elementary particles are tiny (like photons), they are subject to the rules of quantum mechanics. We must therefore abide by our earlier principle: we can only define physical processes by what we can measure in an experiment, and this means we need to be explicit about exactly *what* we can observe about particle interactions. According to the the arguments of the last section, the computation of the probability of any observed process $A \to B$ occurring must involve a sum over all the different ways that process can happen.

[24]This approach to quantum mechanics is known as the sum-over-paths formalism, because it involves summing over all possible paths a particle or system can take. It is due to Richard Feynman who pioneered this approach. Although this view of quantum mechanics is mathematically equivalent to other formulations (such as those originated by Schrödinger, Heisenberg, or Dirac), it leads to a great conceptual simplification in particle physics and other areas.

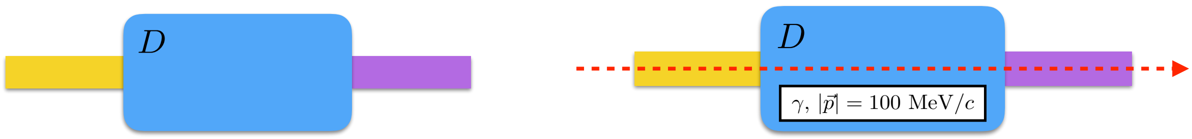

Figure 24: Left pane: An ideal particle detector consisting of a measurement apparatus (blue), intake valve (yellow) and output valve (purple). A particle passing through the detector has its identity determined (*i.e.,* if it's a photon, electron, proton, etc.) and its momentum magnitude measured. This is then output to a screen. Right pane: an illustration of the measurement process, in this case a photon with momentum magnitude $|\vec{p}| = 100$ MeV/c. The direction can be determined by the orientation of the detector.

Let us simplify our lives by considering an ideal particle detector. It is a simple box that tells us the nature of a particle that passes through it (for example, whether it is a proton or electron), as well as the particle's momentum (see Fig. 24). This is illustrated in Fig. 24. We've already seen a simple version of such a detector in the last section where we introduced a box that flashed when a photon passed through it. Now, let's imagine that there is a little LED screen that tells us the name of the particle (photon, electron, proton, etc.) as well as its momentum $\vec{p}$. Of course, in real life we don't have such a magical apparatus, but there do exist detectors that can distinguish different types of particles and measure their momentum, so it is an OK assumption for illustrative purposes[25].

Exercise: Check that MeV/c *is a valid unit of momentum. Recall that electron-Volts is a unit of energy.*

4.2.2 Vertices: The Building Blocks of Particle Interactions

Keeping in the same vein of simplicity, let us consider the most minimal example of a particle physics theory. It is called **quantum electrodynamics (QED)**, and it is the theory of the electron (e^-), its antiparticle the positron (e^+), and the photon. The only force is electromagnetism and it is mediated by the photon. We know that electrons can interact via emitting or absorbing photons, and so the fundamental building block of the theory is the electron-photon interaction. We represent this diagramatically in Fig. 25 and it is known as the **vertex** of the QED theory. The arrows indicate the charge of the particle: in this case, a forward-pointing arrow represents a negatively charged electron, while a backward-pointing arrow represents a positively charged positron. Because the photon is neutral, the conservation of electric charge tells us that there must always be one arrow pointing into the vertex and one arrow pointing out of the vertex so that the net charge sums to zero. This can be in the form of a photon splitting into an electron-positron pair, an electron emitting a photon, or a positron emitting a photon; in each case, it is the same vertex but rotated (all of the rotated versions are shown in Fig. 25). When we change an initial particle to a final particle, we must also change it to its antiparticle: the reason is that it changes a forward arrow to a backward arrow. By convention, we always treat time as pointing *to the right* in representations of particle interactions (this is the modern standard, although you sometimes see in old books that time points from bottom to top).

[25]Of course, in real-life detectors we have to account for the finite resolution of the detector as well as sources of error (for example, the momentum can only be measured to a certain number of decimal places) and we have to contend with particle identification fakes. For example, our detector may indicate that a particle is a photon but it actually was an electron. This complicates matters slightly, but doesn't affect the theoretical formalism too much. For more information on how a particle detector works, see my notes on collider physics.

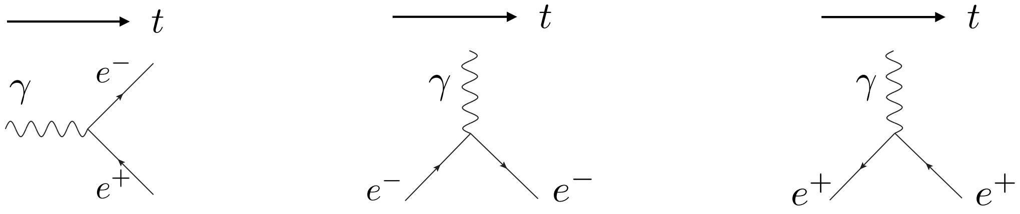

Figure 25: Feynman diagram representation of the interaction between an electron (e^-), positron (e^+), and a photon (γ). Time flows from left to right. Rotating the diagram can represent: (left) a photon splitting into an electron-positron pair; (centre) an electron emitting a photon; (right) a positron emitting a photon.

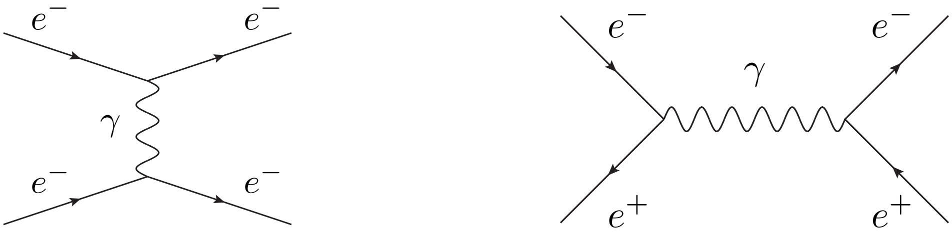

Figure 26: Feynman diagrams showing: (left) the scattering of electrons via the (repulsive) electromagnetic force generated by the exchange of a photon; (right) the annihilation of an electron-positron pair into a photon, which disintegrates back into an electron-positron pair. Remember that time always flows to the right!

Exercise: Can an e^+e^- pair combine to give you a photon? Can two electrons e^-e^- combine to give you a photon? If so, draw the relevant Feynman diagram. If not, explain why not.

4.2.3 Feynman Diagrams: "Paths" For Particle Interactions

Note that we can now assemble different combinations of electrons, positrons, and photons by connecting them with vertices. The result is called a **Feynman diagram**. For example, in the left pane of Fig. 26, we show the scattering of two electrons by exchanging a photon. In the right pane of Fig. 26, we show an electron-positron pair annihilating into two photons. Remember that time flows to the right, and that electric charge conservation requires each vertex to have one arrow flowing in to the vertex and one arrow flowing out.

Ultimately, we want to use these diagrams to characterize physical processes involving elementary particles and their interactions. A common physical process studied in particle physics is *scattering*, where you take two particles, smash them together, and then look at the debris coming out. We measure the incoming particles (their type and momentum) and then the outgoing particles (again, their type and momentum). This is illustrated schematically in Fig. 27.

The crucial part is that we do *not* directly see what happens in the collision. The whole yellow explosion bit in Fig. 27 cannot be directly measured; all we can see is what goes into the collision and what comes out, and then try to reconstruct it from that information[26]. There is therefore a

[26] You may wonder about what happens if we build a more precise detector that is able to peer more deeply into the collision. It's true: we get more information about the collision! However, all that does is re-shuffle part of the yellow collision to the outer components that we are measuring, but there will always be a part we cannot see. The reason is

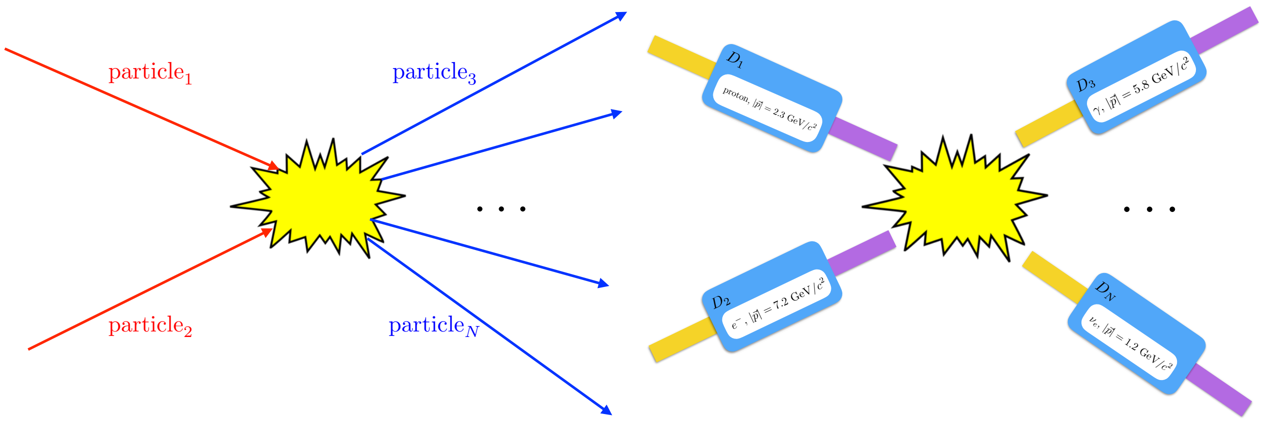

Figure 27: (Left pane) Schematic particle collision: two incoming particles (specified by the type of particle and the momentum of each) collide and transform into $N-2$ final-state particles, again each specified by the type and momentum. (Right pane) We measure what happens in the collision by measuring the types and momenta of each of the initial and final particles (schematic detectors shown), but cannot see directly what happens inside the middle part of the collision.

clear delineation between what is measured in an experiment by our detectors, and what we cannot see.

According to the arguments of Section 4.1, we know that we have to *add* the contributions that are indistinguishable from the point of measurement. Therefore, in calculating the probability of the interaction in Fig. 27, we have to consider every possible force and interaction that allows a proton and an electron to turn into the final-state particles (with the measured momenta). We can make an analogy with the double-slit experiment:

- **Light source:** the analog of the light source is the measured collection of initial-state particles. It is the "known" configuration of the system before the scattering occurs.
- **Screen:** the analog of the screen is the measured collection of final-state particles. We use this information to try to understand what is going on during the collision between the initial pre-collision particles and the final state.
- **Slits:** the analog of the slits (*i.e.,* the different paths of going from the light source to the screen) are the different ways that the Standard Model forces can mediate the collision to go from initial state to final state.

In particle physics, we uncover the properties of particles from their fundamental interactions, and we do so by colliding them, measuring the pre- and post-collision configurations, and then inferring the "paths" or interactions that took them from the initial to final state. We see that particle collisions are simply a more abstract version of the sum-over-paths picture of the double-slit experiment.

Let us consider a concrete example of this last point. Consider again the scattering of an electron with momentum $|\vec{p}_1|$ colliding with an electron with momentum $|\vec{p}_2|$. They turn into an electron of momentum $|\vec{p}_3|$ and an electron of momentum $|\vec{p}_4|$. This is an example of **elastic** scattering, where the mass-energy in the initial and final states is the same, and so the kinetic energy is also the same

that, in quantum mechanics, to see what happens over arbitrarily small distances/short times requires arbitrarily large energies, and so there is always part of the interaction out of reach to us.

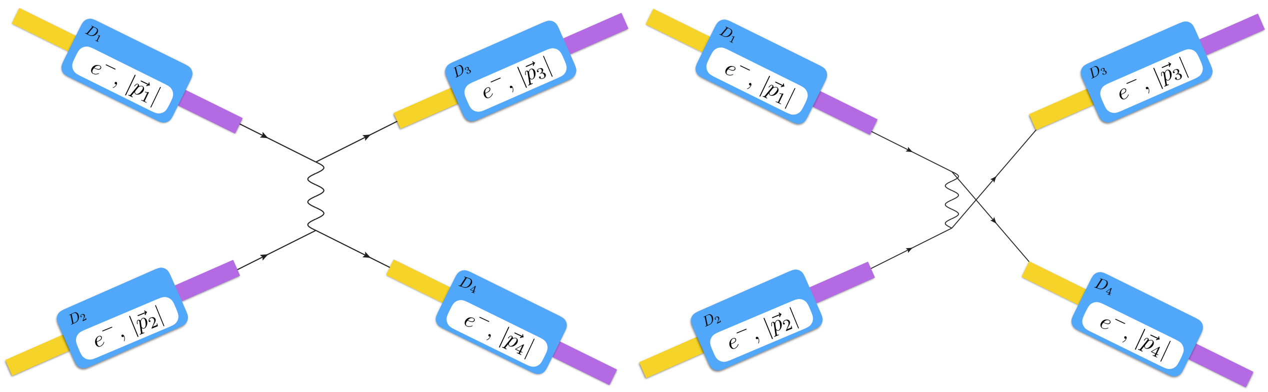

Figure 28: Two incoming electrons collide and transform into two final-state electrons. The two diagrams show the two different ways (or "paths") that map the initial electrons into the final electrons by emitting and absorbing a photon.

between initial and final states: the kinetic energy is simply re-arranged between the electrons in the collision[27]. We drew a Feynman diagram of this process in Fig. 26. But is this the only one?

Let us label the momenta in the Feynman diagram of Fig. 26. The initial electrons have, from top to bottom, momenta $\vec{p}_1$ and $\vec{p}_2$, while the final electrons have, from top to bottom, momenta $\vec{p}_3$ and $\vec{p}_4$. In this diagram, the electron with momentum $\vec{p}_1$ emits a photon and becomes the electron with $\vec{p}_3$, while the photon is absorbed by the electron with momentum $\vec{p}_2$ which becomes the electron with momentum $\vec{p}_4$. This process is illustrated in the left pane of Fig. 28.

However, there is *another* possibility, illustrated in the right pane of Fig. 28. Suppose that the electron with momentum $\vec{p}_1$ emits a photon and then becomes the electron with momentum $\vec{p}_4$. Similarly, the photon is absorbed by electron with momentum $\vec{p}_2$ and becomes the electron with momentum $\vec{p}_3$. At the level of fundamental interactions between electrons and photons, this is a *different* process from the first one considered in the left pane of Fig. 28. However, from the point of view of the detectors, they are physically indistinguishable: the detector D_1 still sees an electron with momentum $\vec{p}_1$ pass through, etc. Therefore, according to the arguments of Section 4.1, we need to *add* the effects of both processes: they are two different "paths" mediating electron-electron scattering.

Putting this all together, what we need to construct the Feynman diagram for a fundamental particle interaction is:

1. Specify the *initial* state(s) of the particle interaction.

2. Specify the *detection* of the final-state particles after the interaction.

3. Using the vertex building block identified in Section 4.2.1, construct all possible diagrams that change the initial state into the final state.

Just like in the double-slit experiment, each diagram (or possible "path") from the initial state to the final state corresponds to a (complex) number, analogous to the wavefunction, related to the probability of the interaction occurring. When there are multiple paths, we have to add together these numbers and take the squared magnitude to get the probability.

[27]This process is known as **Møller scattering**.

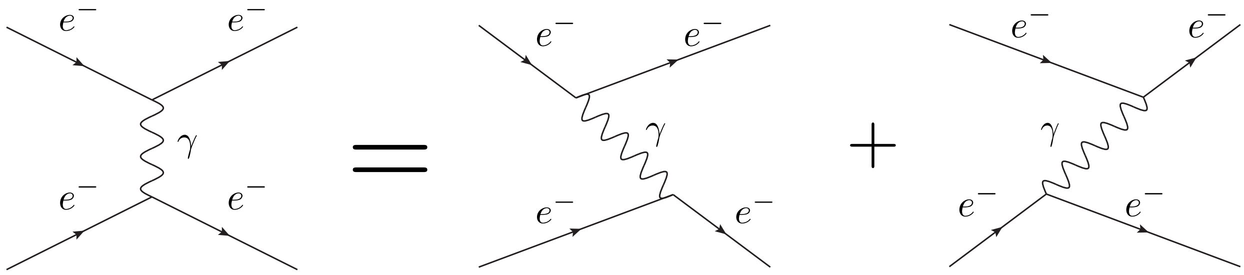

Figure 29: A vertical photon line denotes the combination of: a photon emitted by an electron on the top and absorbed by an electron on the bottom; a photon emitted by an electron on the bottom and absorbed by an electron on the top. (Time always flows to the right in Feynman diagrams.)

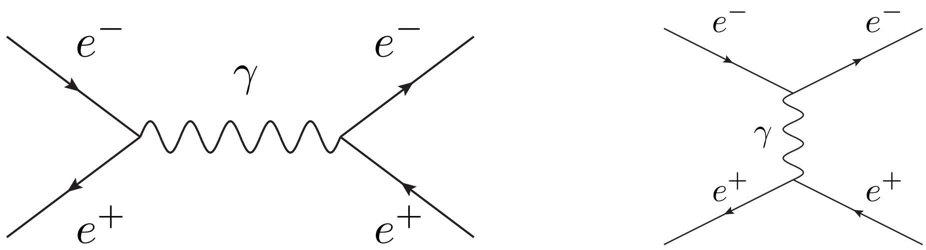

Figure 30: Feynman diagrams for Bhabha scattering $e^-e^+ \to e^-e^+$. The first channel (or path) is an electron-positron annihilating into a photon, which splits back into an e^+e^- pair. In the second channel, one of the electron/positron emits a photon which is absorbed by the other particle.

Before we get to some more examples, there is one more subtlety worth mentioning: in the left pane of Fig. 26, we drew a diagram with a vertical line for the photon. However, you may have expected to have two separate diagrams: one in which the top electron *emits* a photon and the bottom electron *absorbs* the photon, and a diagram in which the reverse happens, namely the bottom electron emits a photon and the top electron absorbs the photon. From the point of view of our detector, we cannot tell these two processes apart. Also, it turns out that they always come in pairs, and so it is simplest to combine them together. Therefore, we do not need to consider separate time orderings of the photon interactions (for instance, whether the top electron absorbs vs. emits the photon as separate processes) because they always are combined into the same process. This is illustrated in Fig. 29.

4.3 More Examples of Feynman Diagrams

We now use what we've learned in the last few sections to draw Feynman diagrams for a couple of processes involving electrons, positrons, and photons. They are:

- **Electron-positron scattering:** This is also known as Bhabha scattering. In this process, we have $e^-e^+ \to e^-e^+$. This can proceed in two ways: either the electron and positron can annihilate to give a photon, which splits back into an e^-e^+ pair, or they can scatter by the exchange of a photon as in $e^-e^- \to e^-e^-$ scattering. These processes are shown in Fig. 30. Note that we do not have a "crossed" Feynman diagram like in the right pane of Fig. 28. The reason is that the electron-photon vertex conserves electric charge and so an electron can not turn into a positron through the emission or absorption of a photon. We also do not include a separate diagram like Fig. 30: any process where the diagrams are the same by simply re-

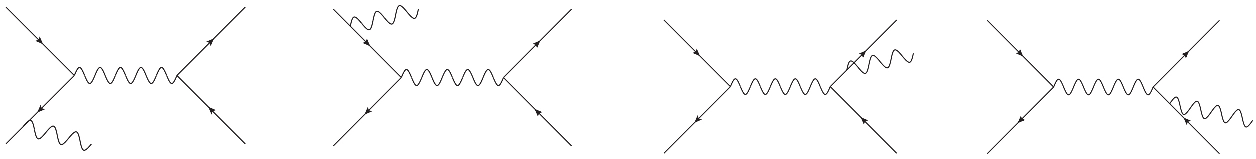

Figure 31: Feynman diagrams showing the reaction $e^-e^+ \to e^-e^+\gamma$. Note that these should also be supplemented with a photon attached to each charged particle in the equivalent of the right pane of Fig. 30.

labelling the lines (as opposed to physically changing which lines correspond to which particles) is the same process and should not be double counted.

- **Photon radiation:** Photons can be radiated from electrons or positrons in collisions via the reaction $e^-e^+ \to e^-e^+\gamma$. To do this, we can take any of the diagrams for electron-positron scattering (for example, the left pane of Fig. 30) and emit a photon from any charged particle. Because we can't tell which charged particle the photon came from, we have to consider each of the diagrams (a subset of the diagrams is shown in Fig. 31).

It takes a lot of practice to be able to identify and draw all of the unique Feynman diagrams contributing to a particular scattering process. At this point, you don't need to have a strong proficiency with this skill, but at least you have the tools to understand what Feynman diagrams mean and how to assemble the vertices or building blocks into diagrams corresponding to different processes.

Exercise: Draw the Feynman diagrams for electron-positron annihilation, $e^-e^+ \to \gamma\gamma$. Your diagrams should include two electromagnetic vertices each. How many different diagrams are there?

Exercise: Draw the Feynman diagrams for light-by-light scattering, $\gamma\gamma \to \gamma\gamma$. Your diagrams should include four electromagnetic vertices. If you have studied classical electromagnetism, how does this result differ from the classical expectation from Maxwell's equations?

4.4 Organizing Diagrams: The Perturbation Series

In the last section, we argued that we must include all Feynman diagrams for processes that are *physically indistinguishable* from one another. Once we specify a list of initial particles (and their momenta) and a list of final particles (and their momenta), we need to draw all diagrams that connect the initial to final particles. For example, in Fig. 28 we saw that we needed *two* diagrams to describe the process $e^-e^- \to e^-e^-$ corresponding to the two cases where a given initial electron directly connects to each of the final electrons.

However, we've cheated a bit. The reason is that, using our electron-photon vertex, we can actually construct more complicated diagrams that *also* contribute to $e^-e^- \to e^-e^-$ scattering. See, for example, the diagrams in Fig. 32. If you look at the initial-state and final-state particles, you see that both of these indeed contribute to $e^-e^- \to e^-e^-$. However, they have more electromagnetic vertices (or building blocks) in them. For example, in the left pane the photon that is emitted from the top electron splits into an electron-positron pair, which then re-annihilate into a photon that continues on its way to be absorbed by the bottom electron. In the right pane, a photon is emitted from the top electron, which then emits a second photon to transmit the force to the bottom electron.

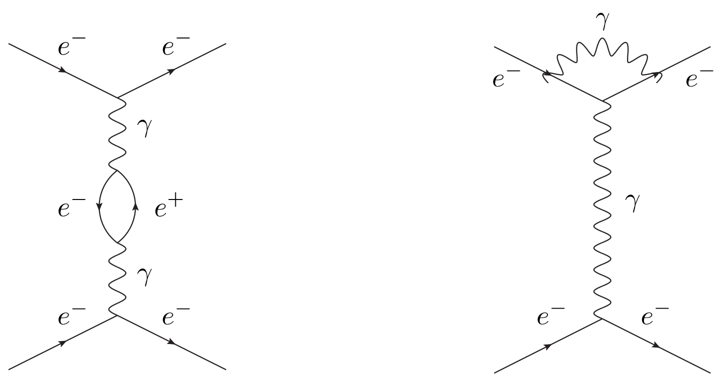

Figure 32: Additional Feynman diagrams contributing to $e^-e^- \to e^-e^-$. These are called loop diagrams due to the close loops of particles appearing in the diagram, as opposed to those in Fig. 28 which are called tree diagrams.

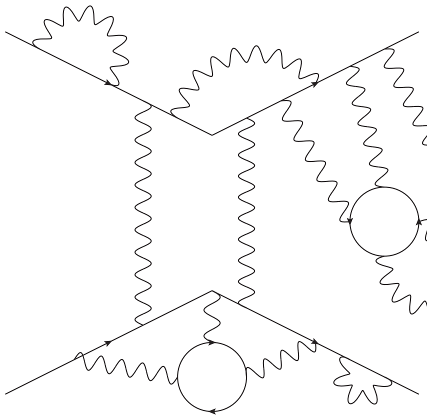

Figure 33: Additional Feynman diagram contributing to $e^-e^- \to e^-e^-$ with many electromagnetic vertices and loops.

The top electron then re-absorbs the first photon it emitted. From the point of view of an external detector, each of these is physically indistinguishable from the diagrams in Fig. 28. They are referred to as **loop diagrams** because of the closed loops of particles in the diagram, as opposed to those in Fig. 28 which are called **tree diagrams**.

At first glance, this seems not to be a problem - it just means we have to include more diagrams than we had originally realized in our characterization of a physical process. However, a little more thought suggests that this is actually a catastrophe! The reason is that there are an *infinite* number of diagrams that correspond to each physical process. For example, in Fig. 33 we show a contribution to $e^-e^- \to e^-e^-$ scattering with 12 loops and 26 electromagnetic vertices. In principle, we have to include this contribution in any calculation of the probability of $e^-e^- \to e^-e^-$ scattering, along with an infinite number of other diagrams! How is that we can make any predictions at all?

A resolution can be found by imagining that the total rate of $e^-e^- \to e^-e^-$ involves summing over an infinite number of paths or diagrams. This rate can only be finite if, as the number of diagrams grows, their contribution to the sum shrinks such that the total sum can still be finite. For example, we can recall that infinite sums like

$$\sum_{k=1}^{N} c_k\, x^k \tag{18}$$

can be finite in the limit $N \to \infty$ for $x < 1$ because, even though there are an infinite number of terms, the terms can become increasingly small as N gets large.

This line of argument turns out to be our saving grace. It turns out that, for every electron-photon vertex in a diagram, the contribution to the rate is proportional to a factor of α, where α is a dimensionless number that characterizes the strength of the force[28]. For the electromagnetic force, we found in Section 2.1 that $\alpha_{\text{EM}} \approx 1/137$.

We can now compare the relative importance of the different types of diagram. Each of the tree diagrams in Fig. 28 has two electromagnetic vertices and so the rate goes like α^2. In contrast, the loop diagrams in Fig. 32 have four vertices and their contribution to the rate goes like α^4, which contains an additional factor of $\alpha^2 \approx 5 \times 10^{-5}$ smaller than the tree diagram. Similarly, the diagram in Fig. 33 contains 26 vertices and so the rate goes like α^{26}, or a factor of $\alpha^{24} \sim 5 \times 10^{-52}$ smaller than the tree contribution. This is good news: it means that, when making predictions, we can *usually* content ourselves with tree-level diagrams. If we are doing precision measurements and comparing the results to theory, we may have to compute higher-loop diagrams. However, we expect that the most complicated diagrams can be ignored in our calculations[29].

The different diagrams of increasing complexity are interpreted as part of a **perturbation series** valid in the $\alpha \to 0$ limit. The idea is that the physical process is dominated by the process with the lowest number of electromagnetic interaction points, and each of the more complicated diagrams gives a small "perturbation" that changes the result by a (hopefully) small result. We can, for example, write the total contribution to $e^-e^- \to e^-e^-$ scattering as

$$\mathcal{P}(e^-e^- \to e^-e^-) = \sum_{k\geq 1} c_k\, \alpha^k. \tag{19}$$

where the first term represents the contribution from the tree diagrams in Fig. 28, the α^4 term represents the contribution from loop diagrams as in Fig. 32, and so forth[30]. The c_k are numerical coefficients that result from calculating the contributions of the various diagrams. Note that in the limit $\alpha \to 0$, there is no interaction whatsoever: the electrons are said to be "free" and do not talk to one another at all.

Exercise: Above, you drew the diagram for light-by-light scattering, $\gamma\gamma \to \gamma\gamma$. In classical electromagnetism, light does not interact with itself because it is not electrically charged, whereas the particle-physics result seems to contradict this fact. Explain why the classical approximation of zero photon-photon interactions works reasonably well even though $\gamma\gamma \to \gamma\gamma$ is a valid diagram.

[28] More accurately, you get a contribution of $\sqrt{\alpha}$ from each vertex to the wavefunction, and then you square the sum of diagrams. But if we consider only a single diagram, then it is equivalent to a power of α for each vertex.

[29] The discussion in this section is overly simplistic for at least one reason: while each diagram does indeed have an additional power of α in the rate for each electromagnetic vertex, the *number* of diagrams tends to grow with the number of vertices. As a result, it is believed that the series is *not*, in fact, convergent. Instead they are asymptotic, which means that they do not converge although truncating the series at a fixed number of terms gives a good approximation to the series for sufficiently small values of α.

[30] Note that terms with odd powers of α can also appear. The reason is that we can also have *interference* between the diagrams, and the rate actually goes like $|d_1\alpha + d_2\alpha^2 + \ldots|^2$, where the first term comes from the tree-diagram processes, the second from the one-loop processes, and so on (d_i are numbers associated with the numerical factors coming from each diagram).

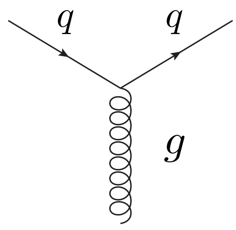

Figure 34: Feynman diagram representation of the interaction between a quark (q) and a gluon (the carrier of the strong force, g).

We can consider Feynman diagrams for the other forces as well. For example, the fundamental vertex of the strong interaction is between a quark and a gluon as shown in Fig. 38, where the gluon is indicated with a curly line. Then, we can construct Feynman diagrams for quark interactions mediated by the strong force by combining this vertex in various ways just as we did with electromagnetism. If we refer to Section 2.3, we see that at the energy of the LEP collider at CERN, we have $\alpha_{\rm s} \approx 0.1$. Therefore, we can still write down a perturbation series where the simpler diagrams are more important and the complicated loop diagrams are suppressed by $\alpha_{\rm s}^n$ for some power n. However, the higher-order α^n pieces are *more* important than for electromagnetism because $\alpha_{\rm s} > \alpha_{\rm EM}$. Indeed, this is what we mean when we say the force is *stronger*: it is more common for interactions to occur via the strong force, and so we have to be more careful and possibly include some of the more complicated loop-diagram processes in our calculations.

There is a complication due to the fact that $\alpha_{\rm s}$ changes with energy[31] as shown in Fig. 13. Indeed, $\alpha_{\rm s}$ becomes larger than 1 for energies $\lesssim 250$ MeV. Once $\alpha_{\rm s} \sim 1$, then the higher-order terms in the series $\alpha_{\rm s}^n$ are no longer suppressed relative to the leading term! In other words, complicated diagrams like Fig. 33 (with the photon lines changed into curly gluon lines) become just as important as the simple tree diagrams. This is a catastrophe: we can no longer rely on simple calculations to predict particle properties and interaction rates. This is known as a **breakdown of perturbation theory** because the more complicated processes become as important (if not important) as the simple processes.

What happens for energies $E \lesssim 250$ MeV? Well, we can no longer even *define* individual quarks and gluons, because each one immediately emits sprays of more quarks and gluons (there is no penalty to do so if $\alpha_{\rm s} > 1$ and these particles cannot be separated). Instead, the strong force becomes like a very powerful glue that binds the quarks and gluons together into *hadrons*, which we have already seen in Section 3.4. We can no longer write down Feynman diagrams with quarks and gluons at such energies, but we can potentially work directly with the resulting protons, neutrons and other hadrons in calculating their interaction rates.

4.5 Interpreting Particles in Feynman Diagrams: Real or Virtual?

Before concluding our section on Feynman diagrams, I want to have a short discussion on what the physical significance is of the particles running inside of the diagram. In the left pane of Fig. 35, I repeat one of the diagrams for $e^-e^- \to e^-e^-$ scattering. We narrate this diagram by saying that the electron with $|\vec{p}_1|$ emits a photon and turns into the electron with momentum $|\vec{p}_3|$, while the photon is absorbed by the electron with momentum $|\vec{p}_2|$ and turns into the electron with momentum $|\vec{p}_4|$. However, all we can really resolve is shown in the right pane of Fig. 35; what happens in the middle

[31] $\alpha_{\rm EM}$ also changes with energy, but not in a dramatic way so we can usually ignore it.

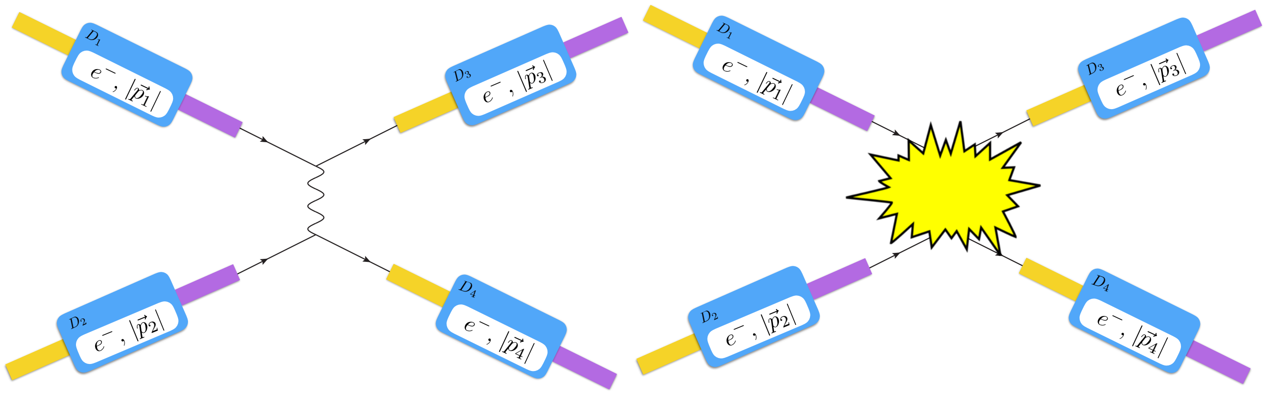

Figure 35: (Left pane) One of the diagrams contributing to $e^-e^- \to e^-e^-$ scattering. (Right pane) The physical scattering process, equivalent to the sum of all intermediate diagrams connecting the initial and final states.

of the "blast" cannot be directly seen. Since we can never directly detect it, is there any sense in which the intermediate photon is "real"?

The answer, in quantum mechanics, is *no*. Only that which we can measure is "real" in any sense that can be unambiguously defined. Indeed, when you get into the details of how you actually calculate the contribution of this diagram to $e^-e^- \to e^-e^-$, you find that the "photon" inside of the diagram has some very un-photon-like properties. For example, a *real* photon in massless, which in the context of special relativity means that it obeys the relation

$$E_\gamma = \sqrt{|\vec{p}_\gamma|^2c^2 + m_\gamma^2c^4} = |\vec{p}_\gamma|c. \tag{20}$$

A particle satisfying this property is said to be **on the mass shell** (or, more succinctly, **on-shell**), which essentially means that the energy and momentum are related to satisfy the relativistic mass-energy relation $m^2c^4 = E^2 - |\vec{p}|^2c^2$. However, it turns out that particles in the middle of Feynman diagrams *do not* need to satisfy this relation (and most often, do not). We say that such particles are **off the mass shell**, or are **off-shell**. For the photon in Fig. 35, you would find that $E_\gamma \neq |\vec{p}_\gamma|c$ for any values of $\vec{p}_1$ and $\vec{p}_3$ (except for the trivial case $\vec{p}_1 = \vec{p}_3$, which corresponds to no interaction in any case).

This seems like a bit of a crisis. Is there a secret group of photons with mass $m_\gamma \neq 0$ that exist in nature that exist only fleetingly in the insides of Feynman diagrams, but are never actually observed? The answer again is *no*. Because we can never actually "see" these particles in a measurement process, there is no sense in which they physically exist. Instead, these intermediate particles are relatively convenient fictions that allow us to describe the workings of forces in particle physics by referring to pointlike particle interactions as shown in Feynman diagrams. They are totally harmless so long as we understand that the inner guts of Feynman diagrams, which can never be directly probed by experiments, should not be taken *too* seriously as real particles. Instead, we refer to them as **virtual particles**[32].

[32]There is a formulation of quantum field theory in which all particles (whether on the inside of the diagram or not) obey the relativistic energy-momentum relationship (or, equivalently, are on-shell). However, this formulation necessitates introducing a violation of conservation of energy in the interactions involving internal particles, even though energy is conserved overall in the collision. To me, this is worse than accepting that internal or virtual particles don't obey the relativistic energy-momentum relationship.

If you look in the literature of particle physics (and, especially, the popular science literature of particle physics), you will see all kinds of statements about how these virtual particles have all kinds of physical effects and can wreak havoc on the laws of physics: for instance, I have seen claims that virtual particles can go backwards in time. Ultimately, these kind of statements are confusing and have no physical meaning: "real" particles that we can see in experiments (*i.e.,* the external legs on Feynman diagrams) have no such wacky properties such as time travel. That is not to say that virtual particles do not obey some physical laws: their interactions conserve energy, momentum, and electric charge. However, because virtual particles are mere useful mathematical tools rather than literal physical objects, we can sidestep any unnerving questions about their physical properties. Indeed, it is possible to formulate physically observable quantities in quantum mechanics without ever referring to virtual particles. We will stick with using virtual particles (as do most particle physicists) because they allow us to describe particle interactions in a convenient way.

Because we are not in a position to delve into the full quantum field theory of particle physics here, we will content ourselves to thinking of Feynman diagrams as point particles mediating interactions by colliding at various spacetime points, keeping in mind not to take the virtual particles inside of Feynman diagrams too seriously. If you would like to learn more about this, you should talk to your research mentor about delving further into quantum mechanics and, eventually, quantum field theory.

5 Feynman Diagrams for the Standard Model Forces

In Section 4, we began our study of how we think about forces and interactions in particle physics. In that case, we focused on a very particular scenario (quantum electrodynamics, consisting of an electron, positron, and photon), and learned how to construct and interpret Feynman diagram representations of particle interactions. As we are well aware by now, there are many more forces than just elecromagnetism, and many more particles than just the electron and positron. In this section, we will go over the interactions for each of the fundamental forces and see how they can be used to construct possible interactions between Standard Model particles.

5.1 The Electromagnetic Force

We start with the electromagnetic force, which is the simplest and most familiar to us. The mediator of the force is the photon, and it is massless and neutral (such that the photon does not directly interact with itself). The photon interacts with any particle that is electrically charged (common jargon is to say that the photon **couples** to charged particles). We can therefore consider any charged Standard Model matter particle X, be it a muon, electron, or up quark. The basic building block for the electromagnetic force is the vertex connecting the photon to the charged particle X and its antiparticle $\bar{X}$. This vertex is shown in Fig. 36.

Exercise: Draw all of the leading-order Feynman diagrams for the following processes: (a) $uu \to uu$; (b) $e^+e^- \to \mu^+\mu^-$. How does your answer to part (b) compare with the Feynman diagrams for $e^+e^- \to e^+e^-$?

The strength of the electromagnetic force is given by the dimensionless constant, $\alpha \approx 1/137$, with a factor of $\sqrt{\alpha}$ contributing to the wavefunction from each vertex (which leads to a factor of α after squaring). However, our derivation of the electromagnetic force in Eq. (5) assumed particles with charge of magnitude e. Other particles have different magnitude charges: for example, u and $\bar{u}$ quarks have charges $Q_u = 2/3\, e$. We therefore expect that the electromagnetic force between up

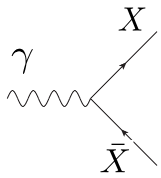

Figure 36: Feynman diagram representation of the interaction between a charged matter particle (X), its antiparticle ($\bar{X}$), and a photon (γ).

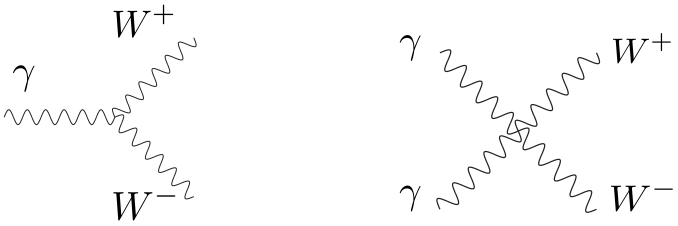

Figure 37: Feynman diagrams showing the basic interaction vertices between $W^{\pm}$ bosons, which have electric charge $\pm e$, and photons. There exist vertices involving a W^+, a W^-, and a photon (left), or a W^+, a W^-, and two photons (right). The magnitude of the wavefunctions representing the probability of each one occurring are still fixed by the single dimensionless coupling, α.

quarks should be *weaker* than between an electron and a positron by a factor of $Q_u^2 = (2/3)^2 = 4/9$, which is less than half of the corresponding rate for electrons. While in Section 4.4 we argued that each vertex between a photon, electron and positron came with a factor of α characterizing the strength of the force in computing the probability of an interaction occurring, we can see that for an arbitrary charged particle X with charge Q_X, the amount that each vertex contributes to the rate of interactions is $Q_X^2\alpha$.

There is one charged particle in the Standard Model that is a force carrier, and not a matter particle. This is the W^+ boson. It has charge $+e$ and therefore feels the electromagnetic force; not surprisingly it interacts with the photon! Because the W^+ boson is a force carrier, we also denote it by a squiggly line. There are two vertices that allow the W^+ and its antiparticle, the W^-, to interact with photon(s). These are illustrated in Fig. 37. Generally, both diagrams will contribute to any process involving photons and $W^{\pm}$ bosons. Note that because the right diagram in Fig. 37 involves a coupling to *two* photons, its rate goes as α^2; you get a power of α in the rate for each time a photon interacts with charged particles.

Exercise: Draw all Feynman diagrams contributing to the process $W^+W^- \to \gamma\gamma$ (there should be three). What is the dependence of the rate on α?

Note that each of the Feynman diagrams for electromagnetic interactions (*i.e.,* interactions of the photon with other particles) have a particular property: they all conserve electric charge! This is because photons **always** produce particles and antiparticles of the same type (for example, an electron and a positron), and since particles and antiparticles have opposite electric charges, the result is always electrically neutral. Similarly, if a particle (such as an electron or a muon or an up quark) emits a photon, it **always** remains the same type of particle. We summarize this by saying that the electromagnetic interactions **conserve electric charge** (if we add up the electric charge

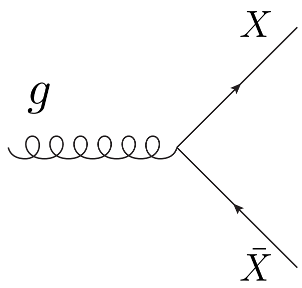

Figure 38: Feynman diagram representation of the interaction between a matter particle charged under the strong interaction (X), its antiparticle ($\bar{X}$), and a gluon (g). The gluon is indicated by a curly line, to distinguish it from the wavy lines of the photon and W/Z bosons.

flowing into a vertex on a Feynman diagram, it always equals the electric charge flowing out of the vertex), and that the electromagnetic interactions **conserve particle type** (they never change an up quark into a charm quark, for example, even though that would conserve electric charge).

5.2 The Strong Force

The strong force is mediated by gluons. Much like the electromagnetic force, the basic building block for Feynman diagrams illustrated strong-force interactions involves a gluon interacting with a particle X and antiparticle $\bar{X}$. The gluon only talks to particles that are charged under the strong force: thus, gluons interact with quarks, antiquarks, and other gluons, but never directly with leptons. In Feynman diagrams, the gluon is indicated by a curly line and interacts with particles that carry colour charge (see Fig. 38).

The fundamental vertex of the strong interactions can be combined to describe more involved scattering processes involving the strong force. For example, in the left pane of Fig. 39 we show a process where an up quark scatters off a charm quark via exchanging a gluon. This process is mediated by the strong force. Because the charm and up quarks are electrically charged, they also experience the same scattering mediated by the electromagnetic force (right pane of Fig. 39); however, because the strong force is much stronger than the electromagnetic force, the largest contribution from scattering occurs due to the diagram with the gluon.

Exercise: The gluon is massless. Can the decay $g \to u\bar{u}$ occur consistent with momentum and energy conservation? Explain.

As discussed in Section 2.3, one of the reasons that the strong force is different from the electromagnetic force is that the gluon is *itself* charged under the strong force. Unlike the photon, which is electrically neutral, the gluon carriers colour charge. This means that the gluon can interact with itself; by contrast, the photon has no direct self-interaction. There are two different vertices by which the gluon can interact with itself. The 3-point and 4-point gluon vertices are shown in Fig. 40. The 3-point vertex gives a contribution to the rate of α_s, while the 4-point vertex gives a contribution to the rate of α_s^2.

Let's consider the process $s\bar{s} \to gg$, illustrated in Fig. 41. The first two diagrams look the same as those for the process $e^+e^- \to \gamma\gamma$. However, the third diagram is found in the strong interactions but not the electromagnetic force: the gluon self-interaction leads to a contribution to quark-antiquark

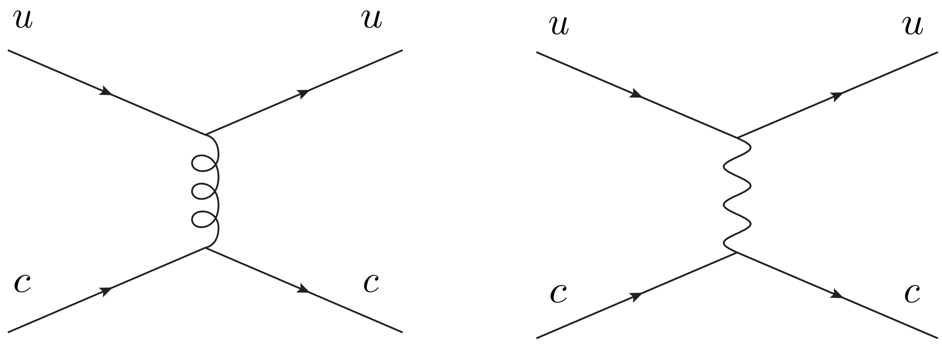

Figure 39: Scattering of an up quark and a charm quark mediated by a (left) gluon; (right) photon. The rate associated with the left diagram goes like α_s^2 while the rate associated with the right diagrams goes like $(2/3)^4\alpha^2$ (the factor of $(2/3)^4$ comes from the fact that the electric charge of the up-type quarks is $2/3e$); since the strong force is much stronger than the electromagnetic force, this means that scattering of quarks is most likely to occur due to gluon exchange, although both processes give some contribution to the overall scattering.

Figure 40: Vertices that show gluon interactions with other gluons. There are two vertices: (left) a 3-point interaction among gluons; (right) a 4-point interaction among gluons.

annihilation.

Exercise: Draw all the Feynman diagrams for $u\bar{u} \to d\bar{d}$.

Exercise: Draw all the Feynman diagrams for $u\bar{u} \to d\bar{d}g$.

Exercise: Draw all the Feynman diagrams for $gg \to gg$ *(there should be four). What is the dependence of the rate of this process on the strong interactions constant,* α_s*? Compare to the photon-photon scattering rate dependence,* $\gamma\gamma \to \gamma\gamma$*, that you found earlier.*

Now you have everything you need to describe the strong interactions! As with the electromagnetic interactions, it is worth pausing and seeing what processes are allowed by the strong interactions and what processes are not allowed. Just like the electromagnetic force, the strong interactions **conserve particle type**. This means, for example, that a vertex between u, $\bar{u}$, and g is allowed, whereas a vertex between u, $\bar{c}$, and g is not allowed. The strong interaction vertices therefore also conserve electric charge (because if the gluon interactions always preserve particle type, electric charge is automatically conserved). Another less obvious property of the strong force is that, by analogy with the electromagnetic force, there exists a **strong charge** that is conserved. It is harder to visualize

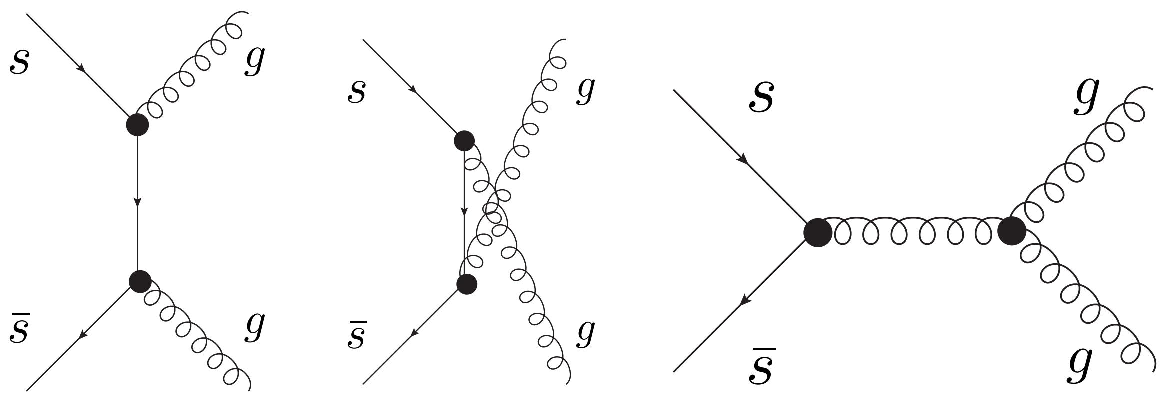

Figure 41: Feynman diagrams for the process $s\bar{s} \to gg$. Each diagram contributes to the rate at order α_s^2, although the diagram on the far right includes a gluon self-interaction contribution to the scattering whereas the first two only include the direct interaction of the strange quark with the gluon.

the strong charge because it is a more complicated mathematical object than electric charge (electric charge is a real number, whereas strong charge is actually represented by a matrix). Leptons don't feel the strong force, and therefore have zero strong charge; quarks and gluons do carry strong charge.

5.3 The Weak Force

The weak force is special in that there are three mediator particles: the Z, W^-, and W^+ bosons. Because the W^+ and W^- are antiparticles of one another, specifying the interactions of one is sufficient to give us the interactions of the other. We now examine each in turn.

Z Boson Interactions

The Z boson is electrically neutral and is its own antiparticle. It is therefore like a heavy version of a photon; however, unlike a photon the Z boson interacts with particle-antiparticle pairs of *every* matter particle, whereas the photon only interacts with those possessing electric charge. As a result, the Z boson interacts with every charged particle *and* with the neutrinos. We show this vertex in Fig. 42; the particle-antiparticle pair labelled by X can be any matter particle in the SM: a quark, charged lepton, or neutrino. As before, this vertex can be rotated to give different processes like $X\bar{X} \to Z$, $X \to ZX$, and $\bar{X} \to Z\bar{X}$.

The Z boson has another property in common with the photon and gluon: it only interacts with particle-antiparticle pairs *of the same type or flavour.* So, for instance, we can have $Z \to \nu_e\bar{\nu}_e$ but not $Z \to \nu_e\bar{\nu}_\mu$.

W Boson Interactions

The interactions of the $W^\pm$ are unique among the SM for one reason: they are electrically charged. Therefore, if we were to have a $W^\pm$ split into a particle-antiparticle pair of the same type, this interaction would violate electric charge: something like $W^+ \to e^+e^-$ or $W^- \to \nu_e\bar{\nu}_e$ is not allowed. Therefore, the $W^\pm$ interactions link a particle and antiparticle *of different types.*

If we consider an interaction $W^- \to X\bar{Y}$, then conservation of electric charge means that the charges of X and $\bar{Y}$ should sum to $-e$, $Q_X + Q_{\bar{Y}} = -e$. Because antiparticles have a flipped sign on

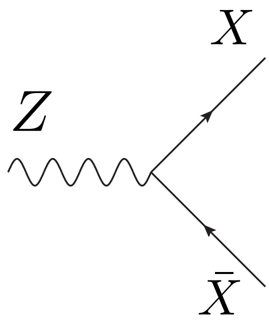

Figure 42: Vertex showing Z boson interaction with a SM particle-antiparticle pair of type X. X can be any matter particle in the SM: the quarks, charged leptons, or neutrinos. However, the Z boson always interacts with particle-antiparticle pairs of the same type, or flavour.

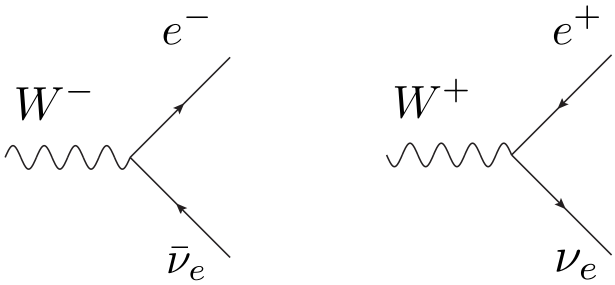

Figure 43: Vertices for the $W^\pm$ interactions with leptons. The W^- couples a charged lepton with an antineutrino *of the same generation*. The cross-generation couplings are not allowed. The W^+ interactions are found by replacing all particles with antiparticles.

the electric charge compared to particles, $Q_{\bar{Y}} = -Q_Y$ and hence $Q_X - Q_Y = -e$. We are therefore looking for the W^- to couple to a particle-antiparticle pair where the charges differ by $-e$.

If you go back to Sec. 3.3, you'll see this is precisely the difference in charge between the down-type and up-type quarks. This is also the difference in charge between the charged leptons and neutrinos. This is not a coincidence: the electric charges of SM particles work out to allow these types of particles to interact with the W bosons.

Let's start with the W^- coupling to leptons. There are three possible vertices: $W^- \to e^-\bar{\nu}_e$, $W^- \to \mu^-\bar{\nu}_\mu$, and $W^- \to \tau^-\bar{\nu}_\tau$. Note that while each of these couplings involves two different types of matter (anti)particles, the W^- always interacts with a charged lepton-antineutrino pair *in the same generation*. Thus, we can have $W^- \to e^-\bar{\nu}_e$ but not $W^- \to e^-\bar{\nu}_\mu$. The W^+ interactions are found by switching all particles to antiparticles: for example, $W^+ \to e^+\nu_e$. The vertices coupling the $W^\pm$ to the leptons are shown in Fig. 43. As before, we can get rotate the vertices to get processes like $e^-\bar{\nu}_e \to W^-$, $e^- \to W^-\nu_e$, etc.

We now move on to the W^- coupling to quarks. From electric charge conservation, we can similarly infer that vertices exist like $W^- \to d\bar{u}$ and $W^- \to s\bar{c}$. However, *unlike* leptons, there also exist cross-generation couplings like $W^- \to s\bar{u}$ and $W^- \to b\bar{c}$. As a result, there are 9 vertices with W^- coupling to one of $d/s/b$ and one of $\bar{u}/\bar{c}/\bar{t}$. The W^+ couplings to quarks are the same as for the W^-, except we replace all particles with antiparticles and vice-versa. The vertices are shown in Fig. 44. The couplings of the W^- among all quark-antiquark pairs are not equal: the

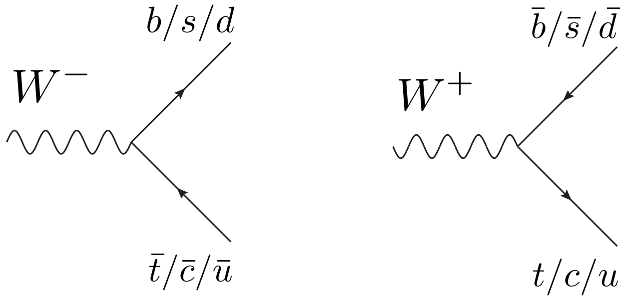

Figure 44: Vertices for the $W^\pm$ interactions with quarks. The W^- couples a charged lepton with an antineutrino *of the same generation.* Cross-generation couplings are allowed, although the couplings are largest within the same generation. The W^+ interactions are found by replacing all particles with antiparticles.

W^- couples preferentially to a quark-antiquark pair within the same generation, while the couplings across generations are suppressed. The larger interaction rate within a generation, and the suppressed coupling across different generations, is illustrated in Fig. 45.

Why is it that the $W^\pm$ interactions are confined to the same generation for leptons, but act across generations for quarks? The answer comes down to the fact that (in the SM) the neutrinos are massless, whereas all of the quarks have non-zero mass. Since all the neutrinos have the same mass (zero), the only thing that distinguishes the different types of neutrinos is what type of charged lepton it interacts with via the $W^\pm$ bosons. For quarks, however, we can distinguish the types of quarks by their mass, and this generically leads to cross-generation couplings of the $W^\pm$ bosons. We don't understand why the couplings within the same quark generation are largest and the cross-generation couplings are suppressed; this is an open question that many particle physicists investigate.

Exercise: Draw the Feynman diagrams for the decay processes $b \to ce^-\bar{\nu}_e$ and $b \to ue^-\bar{\nu}_e$. Which decay do you expect to happen more quickly?.

Finally, there exists a set of vertices between the $W^\pm$ bosons, the Z boson, and the photon. In Sec. 5.1, we saw vertices for $\gamma \leftrightarrow W^+W^-$ and $\gamma\gamma \leftrightarrow W^+W^-$. These are supplemented with purely weak vertices $Z \leftrightarrow W^+W^-$ and $ZZ \leftrightarrow W^+W^-$, as well as a cross-weak-EM vertex $Z\gamma \leftrightarrow W^+W-$ (see Fig. 46). As always, these vertices can be rotated, replacing particles with antiparticles when changing an initial particle to a final particle (or vice-versa). Note that there are no vertices that involve only the photon and Z boson since the Z boson is electrically neutral and does not interact directly with the photon.

5.4 Conservation Laws

If you've studied mechanics, you will have learned several conservation laws, such as conservation of energy, momentum, and angular momentum. In chemistry, you may also have learned about the law of conservation of mass, although this is *not* true in theories that include relativity. These are examples of what we might call **kinematic conservation laws**: assuming some initial state before a collision, they impose restrictions on the momenta and energies after the collision. These conservation laws are certainly important, but they are not the focus of this section.

Instead, we are interested in **dynamical conservation laws**, namely what quantities are conserved by the nature of the force law, rather than restrictions from relativity. For example, looking

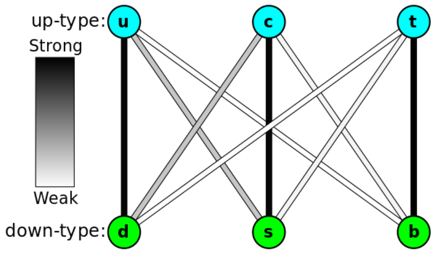

Figure 45: Illustration of the strength of $W^{\pm}$ couplings to particles and antiparticles of different generations of up-type and down-type quarks (image source: Wikipedia, based on C. Amsler *et al.*, *Phys. Lett. B*667 (2008).)

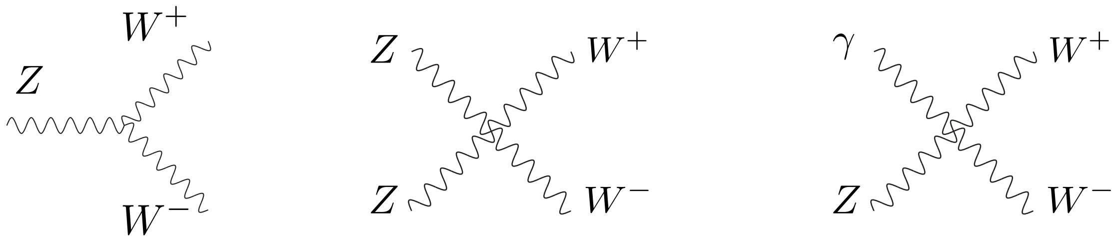

Figure 46: Vertices for interactions involving the $W^{\pm}$ bosons, Z boson, and photon.

strictly at conservation of energy and momentum, we might conclude that the process $e^-e^- \to e^+e^+$ can occur. But this violates conservation of electric charge, which is forbidden in electromagnetism (both classical and quantum versions of it). So even though something is consistent with kinematic conservation laws, they might be inconsistent with conservation laws imposed by the forces in nature. One way to distinguish between the types of conservation laws is as follows: if a process is forbidden from occurring *because I can't draw a Feynman diagram that allows it to happen*, then it's a dynamical conservation laws. If, however, I can draw a Feynman diagram but can't allow the reaction to occur because it's inconsistent with either momentum or energy conservation (an example would be a lighter particle, such as an electron, "decaying" into a heavier particle, such as a muon) then it's a kinematic conservation law at work.

The Standard Model forces encode absolute, exact conservation laws. There are also some properties of particles that are conserved by some forces, but not others. The ones we will examine are the following:

1. Electric charge conservation (exact)

2. Strong charge conservation (exact)

3. Baryon number conservation (exact for our purposes)
4. Lepton number conservation (exact for our purposes)
5. Flavour/generation conservation (approximate)

Electric charge conservation

If you go through all of the interactions so far, you'll find that every single one of them conserves electric charge. This is true not only between initial and final states, but at every vertex in a Feynman diagram. As far as we know, electric charge is always conserved, and a violation of this conservation law has never been conserved. This means that any time a process would change electric charge (for example, a proton decaying to a neutrino and a photon), we automatically know that it can never happen according to the Standard Model.

Strong charge conservation

We mentioned above, in our section on the strong force, that interactions with gluons conserve strong (or colour) charge. Just like electric charge conservation, it turns out that this is true for *all* forces. For example, consider the annihilation of an up quark and an anti-up quark into an electron and a positron, $u\bar{u} \to e^+e^-$. This can proceed through either a photon or a Z boson, neither of which carries strong charge. This means that the colour of the u quark must *match* the anti-colour of the $\bar{u}$ antiquark. For example, if the up quark carries "red" colour, then the $\bar{u}$ carries "anti-red" colour.

The conservation of strong charge is somewhat complicated by the fact that there are three charges, and there is an interplay between them. When we first encountered baryons (such as protons), we found that they are made of 3 quarks, each of which has a different colour charge. However, the *sum* of the three different colour charges is equivalent to no strong charge at all! One way of writing this particular feature is

$$\text{red} + \text{blue} + \text{green} = \text{neutral}. \tag{21}$$

However, we also know that

$$\text{red} + \text{antired} = \text{neutral}. \tag{22}$$

We are forced to conclude that a blue charge plus a green charge is *equivalent* to an antired charge! This means that conservation of strong charge allows *in principle* some processes to occur like

$$u(\text{blue}) + d(\text{green}) \to \bar{d}(\text{antired}) + \gamma. \tag{23}$$

You can check for yourself that this process conserves electric charge, and as we have just argued, it also conserves strong charge. We will soon see that this process violates a different conservation law, however, and so it doesn't actually happen in the Standard Model.

Baryon number conservation

We have just seen that conservation of strong charge is somewhat subtle. In fact, conservation of strong charge and electromagnetic charge alone suggest that certain processes like

$$p^+ \to \pi^+\pi^0 \tag{24}$$

can take place, where p^+ is a proton made of uud quarks, π^+ is a charged pion made of $u\bar{d}$ quarks, and π^0 is a neutral pion made of $u\bar{u}$ and/or $d\bar{d}$ quarks. Both the initial state and final state carry no net strong charge and $+e$ electric charge, and so it is not forbidden by either of our last two conservation laws.

If this process could occur, though, it would lead to disaster: the protons that make up the core of nuclei could decay, and atoms would simply disintegrate. Life as we know it would be completely impossible! It is also in blatant contradiction to observations that protons seem to have persisted for the age of the Universe. Indeed, the current strongest limits on proton decay are that the decay lifetime of a proton has to be longer than about 10^{34} seconds!

It turns that the reason the proton is so long lived is that there exists in the Standard Model an additional conservation law called **baryon number conservation**. Recall that a "baryon" is a fancy name for a proton or neutron, a particle consisting of 3 quarks. The law of baryon number conservation says that a baryon can't simply disappear: once I have a baryon, it always has to decay into another baryon. A baryon can be eliminated by annihilating with an antibaryon: we typically assign baryons a "charge" of $+1$ and antibaryons a "charge" of -1, so that a baryon and antibaryon can annihilate into a non-baryonic state, but otherwise we're stuck with the baryons we have. Since the proton is the lowest-mass baryon, the Standard Model automatically predicts that it is stable.

Since we know deep down that baryons are actually made of quarks, it is conventional to assign quarks a baryon number of $+1/3$ and antiquarks a baryon number of $-1/3$. Therefore, a proton has baryon number $3(+1/3) = +1$, while a meson (like a pion) which is made of a quark + antiquark has baryon number $+1/3 + (-1/3) = 0$. You should check via explicit calculation that the law of baryon number conservation forbids processes like $p^+ \to \pi^+\pi^0$ and $ud \to \bar{d}\gamma$. If you want, you can also take the vertices of the Standard Model forces that we've seen so far and try to draw a Feynman diagram that would permit these processes: try as hard as you might, you will find that simply can't assemble the pieces in such a way as to permit these interactions to happen!

There are a couple of additional comments I want to make. One thing that distinguishes the conservation of baryon number from the conservation of electric or strong charge is that, as far as we know, there is no force that couples to baryon number. In other words, electric charge tells us something about the strength of the coupling between electrically charged particles and photons, but baryon number tells us no such thing. This puts baryon number on a different sort of footing. Indeed, it is generally much easier to find theories that violate conservation laws like baryon number which don't couple to forces, whereas if the charge of a force law is itself violated this can lead to serious problems in the consistency of the theory.

Another comment that is worth making is that, while baryon number is conserved in the Standard Model, it is often *not* conserved in extensions of the Standard Model. There exist a class of theories known as **grand unified theories** which try to unify the strong, weak, and electromagnetic forces into a single super-force. These theories all violate baryon number and predict proton decay (albeit at a very slow rate). In other words, once we add new particles and forces to the Standard Model, we also add to the list of possible Feynman diagrams we can draw, and some of these might lead to baryon number violation. The violation of baryon number plays a central role in theories that try to explain the discrepancy in the observed amounts of matter and antimatter in the Universe, so this is of great importance in our understanding of the laws of nature!

Lepton number conservation

Lepton number conservation is very similar to baryon number conservation, but involving the leptons instead of quarks. More concretely, all leptons are assigned lepton number $+1$ and antileptons are assigned lepton number -1. Because the electromagnetic and weak forces always produce a lepton along with an antilepton, the net lepton number never changes. An example of a lepton-number-violating interaction would be $Z \to \nu_e \nu_e$, but this doesn't actually occur in nature: the only allowed vertex is for $Z \to \nu_e \bar{\nu}_e$. When leptons decay, they still conserve lepton number: for example, a muon decays according to $\mu^- \to e^- \bar{\nu}_e \nu_\mu$, which we can see has lepton number of +1 both before and after the decay.

Flavour/generation conservation

When enumerating the particles of the SM, we saw that each type of particle came in three varieties (or generations, or flavours). For example, there exists an electron, as well as its heavier siblings the muon and tau. The up quark is the lightest flavour of up-type quark, but then there are charm and top quarks too. The strong and electromagnetic forces conserve flavour (or generation): they will never change one type of charged lepton to another. Thus, a muon can emit a photon and stay a muon ($\mu^- \to \gamma \mu^-$), but could not turn into an electron in the process ($\mu^- \to \gamma e^-$). This leads to an important principle: **an elementary particle cannot decay through either the strong or electromagnetic force**. The reason is that these forces conserve particle type, and so if you start with a top quark, you can emit as many gluons and photons as you like and you will always end up with a top quark. Heavier particles cannot decay to lighter ones through the strong and electromagnetic interactions.

The weak force is somewhat more complicated. It turns out that the interactions of the Z boson share the same property as the photon and gluon: they never change particle type. As a result, Z bosons also cannot allow particles to decay in the SM.

The $W^\pm$ boson, however, is the only charged force carrier and that makes it special. The $W^\pm$ boson links up pairs of particles (the up- and down-type quarks, and the charged and neutral leptons) allowing one to turn into the other. For example, the SM allows the top quark to decay via $t \to W^+ b$. Thus, $W^\pm$ boson interactions change particle type.

For leptons, the $W^\pm$ always links the charged and neutral leptons of the same generation. For example, a decay $W^- \to e^- \bar{\nu}_e$ is allowed, while $W^- \to \mu^- \bar{\nu}_e$ is not. We say that there exists a **lepton flavour conservation** law. If we look again at muon decay, $\mu^- \to e^- \bar{\nu}_e \nu_\mu$, we see that we start with a muon lepton number of +1 and an electron lepton number of 0, and in the final state we have the same thing (the e^- and $\bar{\nu}_e$ have opposite electron lepton numbers, whereas ν_μ has muon lepton number +1. For leptons, the weak force therefore conserves flavour or generation number: a τ^+ can decay into a $\bar{\nu}_\tau$ plus other things, but there always has to be a tau-type antilepton in the final state[33].

This is not true for the quarks. While it is *usually* true that a top quark decays via $t \to W^+ b$ since the top and bottom quarks are in the same generation, it still rarely occurs that a top quark can decay via $t \to W^+ s$ os $t \to W^+ d$. In the case of quarks, flavour or generation is *not* conserved by the weak force. This allows top quarks to decay into bottom quarks, bottom quarks into charm quarks, charm quarks into strange quarks, and finally strange quarks into up quarks. The result is that we have nothing hanging around of any of the heavier generations of quarks, whereas in leptons we have neutrinos of all varieties in the present Universe.

The reason why we say that flavour or generation is approximately conserved is that the electromagnetic and strong interactions are so much more powerful at the low energies at which most SM particles are decaying, whereas the weak force processes are much slower. This means that particle decays that conserve flavour happen very fast, whereas particle decays that change flavour happen slower. This hierarchy in decay rates led to the notion of approximate conservation.

[33] Once we take into account neutrino masses, it is possible for neutrinos of one type to turn into another. However, these effects are very small except in dedicated neutrino experiments that are meant to look for them.

A Appendix: Recap of Modern Physics Concepts

A.1 Quantum Physics

The basic idea of quantum physics is that certain systems can only take on particular discrete, "quantized" energy levels. For example, in the classical theory of electromagnetism, an electric or magnetic field can have any intensity, but in quantum mechanics we find that the electromagnetic field is made up of irreducibly tiny lumps of energy "quanta" known as photons. The energy in each photon is determined by the *frequency* (or colour) of the electromagnetic wave, with a ratio given by **Planck's constant**, $h \approx 6.62 \times 10^{-34}$ J $\cdot$ s,

$$E_{\text{photon}} = h\nu, \tag{25}$$

where ν is the frequency in Hz. What this quantization of electromagnetic energy means is that, if I have a beam of light with some frequency ν, there is a *minimum* amount of energy corresponding to a single photon of that frequency of light. This is much different from a classical view of light, where we can make a light sources as dim as we want regardless of the frequency of light.

The concept of wavelength may be more familiar to you than frequency. If you picture a wave consisting of peaks and troughs, the wavelength, λ, is the physical distance between two adjacent peaks. For light travelling at speed c, the relationship between speed, wavelength, and frequency, is $c = \lambda\nu$. Thus, if I have light of a particular wavelength, the corresponding energy of a photon is

$$E = \frac{hc}{\lambda}. \tag{26}$$

The energy in a single photon is quite tiny. If we consider a wavelength corresponding to green light, $\lambda = 550$ nm, the photon has energy 4×10^{-19} J: a truly tiny energy (although it is not so tiny when you consider that it is being carried by a single particle)! Importantly, a particle's energy is *inversely proportional* to its wavelength; this will have important implications when we think about the energies needed to look at small length scales.

If you're confused about how a photon can be both a lump (particle-like) and have a wavelength (wave-like), you are in good company. This is something that is hard for most people to wrap their brains around, and the basic answer is that quantum mechanical objects don't behave in ways that we are used to from classical mechanics. In particular, objects can have two different, contradictory-seeming, properties at the same time! This is something that you will learn to accept over time; for some of you who are used to thinking beyond simple binaries, this may already come easily to you.

In physics, we often like to use angular units, so instead of using Hz (which is cycles per second) we like to use angular frequency (in radians per second). Because one turn of a circle is 2π radians, we have that

$$\omega = 2\pi\nu. \tag{27}$$

To avoid carrying all these extra factors of 2π around, we typically also define a **reduced Planck's constant**, $\hbar = h/2\pi$, and so we have $E = \hbar\omega$. The quantity $\hbar$ is used far more frequently in physics than h itself.

A.1.1 Energy and Spatial Size in Quantum Mechanics

Eq. (26) gives a reciprocal relationship between energy and a length scale, in this case λ. In general, if we want to study objects that are smaller and smaller in size, we must use higher energy particles as probes to bounce off the object. In this subsection, we'll make a few arguments as to why this

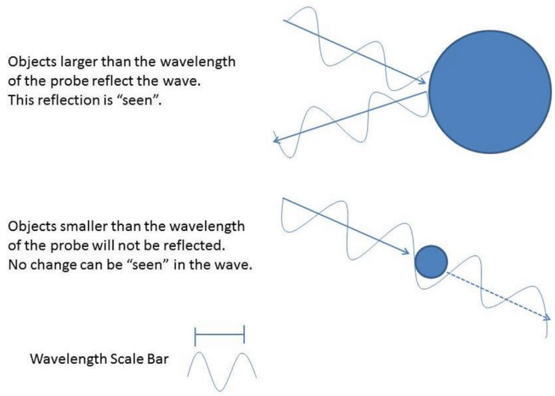

Figure 47: Illustration of different wave behaviours (reflection vs. transmission) depending on relative size of object vs. wavelength (image source: R. E. Michin *et al.*, FDOT Research Project No. BDK75 977-56).

is the case. However, all you really need to know to follow the notes is this reciprocal relationship between energy and size; the rest of this section is optional.

Let's consider looking at an object using visible light. The way we "see" such an object is by looking at the light that is reflected off its surface (or, alternately, by looking at the shadow behind the object). However, whether a wave reflect or diffracts off an object depends on the size of the object compared to the wavelength of light (see Fig. 47). If the wavelength is small compared to the object, then the beam is significantly scattered by it. If, however, the wavelength is much larger than the object, the wave diffracts around it and the object becomes almost undetectable. Since the dimmest beam of light that we can possibly make consists of a single photon and we know that $E \propto 1/\lambda$, this means we need to go to higher energies in order to resolve smaller objects.

In the above example, we considered a classical object with a particular size that we studied by bouncing a quantum wave off of it. Of course, the object we are studying *itself* is quantum in nature, and we now argue that there is a similar reciprocal relationship between an object's energy and its smallest possible spatial extent. In quantum mechanics, a particle no longer has a definite position but it has some uncertainty in its position. The reason is that quantum mechanics is probabilistic in nature, and the state of a particle is given by something called the wavefunction, $\psi(x)$. If the magnitude of $\psi(x_0)$ is larger in one position x_0, then there is a higher probability of detecting the particle there if we go and look for it. If the magnitude of $\psi(x_1)$ is smaller in another location, then there is a low probability of detecting the particle at x_1. The spread in possible positions of the particle is given by the spatial spread in $\psi(x)$. We call such a configuration of a particle a **wave packet**; it's like a little bundle of waves that represents a single particle along with its wave nature.

We provide an example of a wave packet in Fig. 48. In this figure, the wavefunction is illustrated as a function of position. In general, the wavefunction is a complex number; as a result, the real and imaginary parts of $\psi(x)$ are plotted along with the magnitude. It is evident that, in spite of the wiggly nature of the wavefunction, the particle exists in a lump near $x = 1$.

All wave packets have a spatial spread in them: the width of the lump. There are technical ways of defining it, but for our purposes here it is pretty clear that the width of the lump is a little less than 2 units of position. It's fine that there is some non-zero values of $\psi(x)$ far from the lump, but

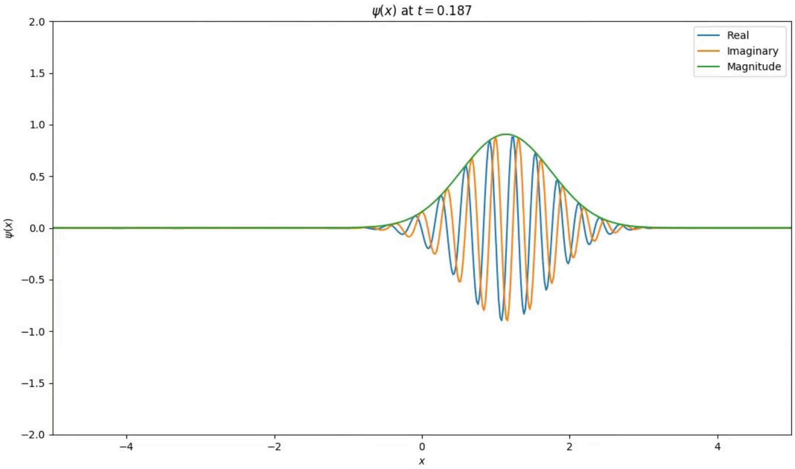

Figure 48: A wave packet showing where the particle is most likely to be found if its position is measured. This is given by the wavefunction, $\psi(x)$: this figure shows the real and imaginary parts of $\psi(x)$, as well as the magnitude $|\psi(x)|$ (image source: D. Roundy, https://www.youtube.com/watch?v=ngGpSyWgDcY).

"most" of the wavefunction should be contained within the width. We define a quantity Δx, which is half the width of the lump. So, in the case of Fig. 48, we have $\Delta x \approx 1$.

In quantum mechanics, it is possible to prove a very deep property known as the **Heisenberg Uncertainty Principle**. Maybe you've heard of it, even if you don't know what it is. Mathematically, the Heisenberg Uncertainty Principle says that if an object has a spatial spread Δx, then it also has a spread in its possible values of momentum, Δp, and that these two quantities are related by

$$\Delta x \, \Delta p \geq \frac{\hbar}{2}. \tag{28}$$

In other words, the narrower I make the spatial spread of my lump, the more uncertain the lump's momentum is. This means if I know *perfectly* well where an object is at one moment in time, then it has a huge spread in momentum and so I don't know where it is at later times.

What does this have to do with energy? For a free particle, the larger the momentum, the larger the kinetic energy. As a result, an uncertainty in momentum leads to an uncertainty in energy. In particular, a particle with a large Δp will have a large spread of energies too! The precise relationship depends on whether this is a relativistic or non-relativistic particle: we are usually interested in relativistic particles, for which the relationship between momentum and energy is $E = cp$. It directly follows that $\Delta E = c\Delta p$, and we have

$$\Delta E \quad \geq \quad \frac{\hbar c}{2\Delta x}. \tag{29}$$

So, the *smaller* the size of an object in space, the larger its spread in energies, and consequently the larger the average energy of the system.

Let's look at some examples. First, let's take an elementary point particle, the electron. Its minimum energy is its mass energy, $E = M_e c^2$. If we want to take an electron that is sitting close to rest and want to know how "large" it is quantum mechanically, then we can plug $E = M_e c^2$ into Eq. (29) and get an estimate of the size,

$$\Delta x \approx \frac{\hbar}{2M_e c} \approx 1.9 \times 10^{-13} \text{ m}. \tag{30}$$

The electron is *really* tiny! This is why we can usually approximate electrons as point particles[34].

For our second example, consider the strong force, which becomes very strong for particles with energies $E \approx 200$ MeV $\approx 3 \times 10^{-11}$ J. Mediators of the strong force, the gluons, therefore typically have this energy to hold quarks together inside of protons and neutrons. Using Eq. (29), The corresponding spread in positions is approximately 10^{-15} m, which is the size of the proton.

Our discussion of energy and spatial size is necessarily somewhat sketchy. All of the relationships between energy and spatial size can be derived more rigorously in quantum mechanics. However, you can get the correct physical picture with some of the ideas explored here.

A.2 Special Relativity

In particle physics, we typically study particles that are moving at speeds comparable to the speed of light. We therefore must operate in the framework of special relativity. In these notes, we largely explore the qualitative features of modern theories of particle physics, and therefore we are most interested in the underlying principles of special relativity rather than calculational techniques. For more details on these, please see my notes on collider physics.

The foundational concept of special relativity is that there exists a maximum velocity that is universal to all observers, namely the speed of light (denoted c). More importantly, the speed of light is the same to every observer; this means that someone on a rocket ship passing rapidly by Earth will observe light from the Sun to move at the same velocity as we would on Earth. This leads to all sorts of important conclusions, such as the lack of a universal definition of time, the fact that events that are simultaneous for one observer are not simultaneous for another obsever, and so on.

Of most relevance to particle physics is that, because all observers will agree on the speed of an object travelling at speed c, there is no observer or frame of reference in which that object is at rest. In other words, that object simply cannot exist at rest; we can never see "stopped" light[35]. By contrast, any object that is moving *slower* than the speed of light will have different velocities for different observers. In particular, there will be one observer that sees the object at rest. This is consequently known as the **rest frame of reference**. You are familiar with this effect, for instance,

[34]For those who have studied some chemistry and/or quantum mechanics, you will know that the typical size of an orbital in a hydrogen atom is the Bohr radius, $a_0 \sim 10^{-10}$ m. If you take this as the spread in positions of an electron, it is much larger than what we just calculated! The reason is that for the electromagnetic interactions inside of an atom, an electron can never turn into anything else: it's mass-energy is conserved. Consequently, no process actually changes by an energy that is anywhere close to the electron mass-energy, and consequently its spread in positions can be much larger. If, however, we fire a positron at an atom and intend to annihilate the electron, then we have to be as close as 10^{-13} m in order for them to meet and annihilate. Thus, the position uncertainty resulting from the mass-energy relation is the *irreducible* uncertainty associated with the energy in a particle's mass, even if there are additional uncertainties arising, for example, from its atomic orbital. In any case, the relation we have derived is only valid for a free electron, not an electron inside of an atom.

[35]Light can *appear* to travel slower than c when passing through a medium, such as water or glass; however, it is not that light is actually moving slower than c, but that it is bouncing around interacting with atoms in the material and so its *effective* forward speed is reduced.

if you are sitting on a moving train: your rest frame coincides with the train (you aren't moving relative to your seat, for instance) and you see the world moving backwards out the window, while an observer on the ground will see you moving forward.

We can define the **rest energy** of an object as the energy it possesses when it is not moving (*i.e.,* when examined in its rest frame). This is an intrinsic property of the object (determined by, say, its composition, geometry, whatever), and is the minimum amount of energy associated with the existence of the object: an observer can see it having *larger* energy than its rest energy if the object is moving, but it can never have *smaller* energy than its rest energy. In fact, we can use this to define the **mass**, M, of the object,

$$M \equiv \frac{E_{\text{rest}}}{c^2}. \tag{31}$$

This is simply a relation from dimensional analysis: we know that both E_{rest} is an intrinsic property of the object, and we can multiply (or divide) it by a constant in order to obtain an equivalent quantity that has dimensions of mass.

By examining the object in a different frame in which it is moving with momentum $\vec{p}$, the energy is obviously larger: it has a kinetic energy in this frame in addition to its rest (or mass) energy. It is possible to use the Lorentz transformations of special relativity to show that the energy in a general frame is

$$E = \sqrt{|\vec{p}|^2c^2 + M^2c^4}. \tag{32}$$

We can see that, in the $|\vec{p}| \to 0$ limit, we recover $E_{\text{rest}} = Mc^2$, while in the $M \to 0$ limit we obtain $E = |\vec{p}|c$. In this case, we have energy proportional to momentum: for a massless particle, there is no energy in the rest frame. This is consistent with our picture of a particle travelling at the speed of light, which *does not have* a rest frame. Thus, we see that a particle travelling at the speed of light is a massless particle.

Exercise: So far, we have defined a quantity M which we have called the "mass" and related it to the energy of an object in its rest frame. However, we have not yet connected it to the concept of "mass" traditionally taught in introductory courses. By performing a Taylor series expansion on E in the $|\vec{p}| \to 0$ limit, show that the energy reduces to a sum of a "rest energy" piece and the conventional non-relativistic kinetic energy where we interpret M as a mass.

There are a few other important properties of special relativity that will be relevant to us. One is the time dilation effect. The statement of time dilation is that, if we take a time Δt measured in the rest frame of an observer, then the time measured in a frame moving at constant speed v relative to the rest frame is

$$\Delta t' = \frac{\Delta t}{\sqrt{1 - v^2/c^2}}. \tag{33}$$

Note that this blows up as $v \to c$, another sign that we cannot define a rest frame for an object moving with $v = c$! Most importantly, we see that $\Delta t'$ is always *larger* than Δt. This has physical implications: for example, the typical lifetime of a particle called a pion in its rest frame is 2.6×10^{-8} s, after which time it decays into lighter particles. For a pion moving at approximately the speed of light $c \sim 3 \times 10^8$ m/s, we expect it should travel a distance ~ 10 m, which is the typical size of a detector in a particle physics detector. However, the vast majority of pions produced at the Large Hadron Collider *do not decay* inside of the detector! The reason is that they are moving at a very high speed $v \sim c$ relative to the detector in the laboratory, and so its decay time in the laboratory frame is larger by a factor of $1/\sqrt{1 - v^2/c^2} \gg 1$.

This can quickly get very confusing. The solution is, where possible, to compute physical quantities such as decay times in the rest frame of the object. Then, by applying the appropriate formulae from special relativity, it is straightforward to figure out the relevant times for all other observers.

There is a final topic that I should address: what we are calling "mass" is sometimes called "rest mass". This is to be contrasted with "relativistic mass", which is defined as $M/\sqrt{1 - v^2/c^2}$. The problem with relativistic mass is that it is *different for each observer*, whereas we understand mass in particle physics as a property of elementary particles that is *the same for each observer*: it is related to the fundamental, irreducible amount of energy stored in the particle when it is at rest. Indeed, each particle in every high-energy particle interaction will have a slightly different speed; this would be a real mess if we had to define a separate mass for each! Therefore, we once and for all dispose of the notion of "relativistic mass", and the mass of the particle is something that is the same for all observables.