1 Introduction

1.1 Preface and conventions

This book is meant as a quick and dirty introduction to the techniques of quantum field theory. It was inspired by a little book (long out of print) by F. Mandl, which my advisor gave me to read in my first year of graduate school in 1969. Mandl's book enabled the smart student to master the elements of field theory, as it was known in the early 1960s, in about two intense weeks of self-study. The body of field-theory knowledge has grown way beyond what was known then, and a book with similar intent has to be larger and will take longer to absorb. I hope that what I have written here will fill that Mandl niche: enough coverage to at least touch on most important topics, but short enough to be mastered in a semester or less. The most important omissions will be supersymmetry (which deserves a book of its own) and finite-temperature field theory. Pedagogically, this book can be used in three ways. Chapters 1–6 can be used as a text for a one-semester introductory course, the whole book for a one-year course. In either case, the instructor will want to turn some of the starred exercises into lecture material. Finally, the book was designed for self-study, and can be assigned as a supplementary text. My own opinion is that a complete course in modern quantum field theory needs 3–4 semesters, and should cover supersymmetric and finite-temperature field theory.

This statement of intent has governed the style of the book. I have tried to be terse rather than discursive (my natural default) and, *most importantly, I have left many important points of the development for the exercises. The student should not imagine that he/she can master the material in this book without doing at least those exercises marked with a* *. In addition, at various points in the text I will invite the reader to prove something, or state results without proof. The diligent reader will take these as extra exercises. This book may appear to the student to require more work than do texts that try to spoon-feed the reader. I believe strongly that a lot of the material in quantum field theory can be learned well only by working with your hands. Reading or listening to someone's explanation, no matter how simple, will not make you an adept. My hope is that the hints in the text will be enough to let the student master the exercises and come out of this experience with a thorough mastery of the basics.

The book also has an emphasis on theoretical ideas rather than application to experiment. This is partly due to the fact that there already exist excellent texts that concentrate on experimental applications, partly due to the desire for brevity, and partly to increase the shelf life of the volume. The experiments of today are unlikely to be of intense interest even to experimentalists of a decade hence. The structure of quantum field theory will exist forever.

Throughout the book I use natural units, where $\hbar = c = 1$. Everything has units of some power of mass/energy. High-energy experiments and theory usually concentrate on the energy range between 10^{-3} and 10^3 GeV and I will often use these units. Another convenient unit of energy is the natural one defined by gravitation: the Planck mass, $M_P \approx 10^{19}$ GeV, or reduced Planck mass, $m_P \approx 2 \times 10^{18}$ GeV. The GeV is the natural unit for hadron masses. Around 0.15 GeV is the scale at which strong interactions become strong. Around 250 GeV is the natural scale of electro-weak interactions, and $\sim 2 \times 10^{16}$ GeV appears to be the scale at which electro-weak and strong interactions are unified.

I will use non-relativistic normalization, $\langle p|q\rangle = \delta^3(p-q)$, for single-particle states. Four-vectors will have names which are single Latin letters, while 3-vectors will be written in bold face. I will use Greek mid-alphabet letters for tensor indices, and Latin early-alphabet letters for spinors. Mid-alphabet Latin letters will be 3-vector components. I will stick to the van der Waerden dot convention (Chapter 5) for distinguishing left- and right-handed Weyl spinors. As for the metric on Minkowski space, I will use the West Coast, *mostly minus*, convention of most working particle theorists (and of my toilet training), rather than the East Coast (mostly plus) convention of relativists and string theorists.

Finally, a note about prerequisites. The reader must begin this book with a thorough knowledge of calculus, particularly complex analysis, and a thorough grounding in non-relativistic quantum mechanics, which of course includes expert-level linear algebra. Thorough knowledge of special relativity is also assumed. Detailed knowledge of the mathematical niceties of operator theory is unnecessary. The reader should be familiar with the Einstein summation convention and the totally anti-symmetric Levi-Civita symbol $\epsilon^{a_1 \ldots a_n}$. We use the convention $\epsilon^{0123} = 1$ in Minkowski space. It would be useful to have a prior knowledge of the theory of Lie groups and algebras, at a physics level of rigor, although we will treat some of this material in the text and Appendix G. I have supplied some excellent references [1–4] because this math is crucial to much that we will do. As usual in physics, what is required of your mathematical background is a knowledge of terminology and how to manipulate and calculate, rather than intimate familiarity with rigor and formal proofs.

1.1.1 Acknowledgements

I mostly learned field theory by myself, but I want to thank Nick Wheeler of Reed College for teaching me about path integrals and the beauties of mathematical physics in general. Roman Jackiw deserves credit for handing me Mandl's book, and Carl Bender helped me figure out what an instanton was before the word was invented. Perhaps the most important influence in my grad school years was Steven Weinberg, who taught me his approach to fields and particles, and everything there was to know about broken symmetry. Most of the credit for teaching me things about field theory goes to Lenny Susskind, from whom I learned Wilson's approach to renormalization, lattice gauge theory, and a host of other things throughout my career. Shimon Yankielowicz and Eliezer Rabinovici were my most important collaborators during my years in Israel. We learned a lot of great physics together. During the 1970s, along with everyone else in the field, I learned from the seminal work of D. Gross, S. Coleman, G. 't Hooft, G. Parisi, and E. Witten. Edward was a friend and a major influence throughout my career. As one grows older, it's harder for people to do things that surprise you, but my great friends and sometimes collaborators Michael Dine, Willy Fischler, and Nati Seiberg have constantly done that. Most of the field theory they've taught me goes beyond what is covered in this book. You can find some of it in Michael Dine's recent book from Cambridge University Press.

Field theory can be an abstract subject, but it is physics and it has to be grounded in reality. For me, the most fascinating application of field theory has been to elementary particle physics. My friends Lisa Randall, Yossi Nir, Howie Haber, and, more recently, Scott Thomas have kept me abreast of what's important in the experimental foundation of our field.

In writing this book, I've been helped by M. Dine, H. Haber, J. Mason, L. Motl, A. Shomer, and K. van den Broek, who've read and commented on all or part of the manuscript. The book would look a lot worse than it does without their input. Chapter 10 was included at the behest of A. Strominger, and I thank him for the suggestion. Chris France, Jared Rice, and Lily Yang helped with the figures. Finally, I'd like to thank my wife Ada, who has been patient throughout all the trauma that writing a book like this involves.

1.2 Why quantum field theory?

Students often come into a class in quantum field theory straight out of a course in non-relativistic quantum mechanics. Their natural inclination is to look for a straightforward relativistic generalization of that formalism. A fine place to start would seem to be a covariant classical theory of a single relativistic particle, with space-time position variable $x^{\mu}(\tau)$, written in terms of an arbitrary parametrization τ of the particle's path in space-time.

The first task of a course in field theory is to explain to students why this is not the right way to do things.[1] The argument is straightforward.

Consider a classical machine (an emission source) that has probability amplitude $J_{\mathrm{E}}(x)$ of producing a particle at position x in space-time, and an absorption source, which has amplitude $J_{\mathrm{A}}(x)$ to absorb the particle. Assume that the particle propagates

[1] Then, when they get more sophisticated, you can show them how the particle path formalism can be used, with appropriate care.

Fig. 1.1. **Boosts can reverse causal order for $(x-y)^2 < 0$.**

freely between emission and absorption, and has mass m. The standard rules of quantum mechanics tell us that the amplitude (to leading order in perturbation theory in the sources) for the entire process is (remember our natural units!)

$$A_{\mathrm{AE}} = \int \mathrm{d}^4x\, \mathrm{d}^4y \langle x|\mathrm{e}^{-\mathrm{i}H(x^0-y^0)}|y\rangle J_{\mathrm{A}}(x) J_{\mathrm{E}}(y), \tag{1.1}$$

where $|x\rangle$ is the state of the particle at spatial position x. This doesn't look very Lorentz-covariant. To see whether it is, write the relativistic expression for the energy $H = \sqrt{p^2+m^2} \equiv \omega_p$. Then

$$A_{\mathrm{AE}} = \int \mathrm{d}^4x\, \mathrm{d}^4y\, J_{\mathrm{A}}(x) J_{\mathrm{E}}(y) \int \mathrm{d}^3p |\langle 0|p\rangle|^2 \mathrm{e}^{-\mathrm{i}p(x-y)}. \tag{1.2}$$

The space-time set-up is shown in Figure 1.1. In writing this equation I've used the fact that momentum is the generator of space translations[2] to evaluate position/momentum overlaps in terms of the momentum eigenstate overlap with the state of a particle at the origin. I've also used the fact that (ω_p, p) is a 4-vector to write the exponent as a Lorentz scalar product. So everything is determined by quantum mechanics, translation invariance and the relativistic dispersion relation, up to a function of 3-momentum. We can determine this function up to an overall constant, by insisting that the expression is Lorentz-invariant, if the emission and absorption amplitudes are chosen to transform as scalar functions of space-time. An invariant measure for 4-momentum integration, ensuring that the mass is fixed, is $\mathrm{d}^4p\, \delta(p^2 - m^2)$. Since the momentum is then forced to be time-like, the sign of its time component is also Lorentz-invariant (Problem 2.1). So we can write an invariant measure $\mathrm{d}^4p\, \delta(p^2-m^2)\theta(p^0)$ for positive-energy particles of mass m. On doing the integral over p^0 we find $\mathrm{d}^3p/(2\omega_p)$. Thus, if we choose the normalization

$$\langle 0|p\rangle = \frac{1}{\sqrt{(2\pi)^3 2\omega_p}}, \tag{1.3}$$

then the propagation amplitude will be Lorentz-invariant. The full absorption and emission amplitude will of course depend on the Lorentz frame because of the coordinate dependence of the sources $J_{\mathrm{E,A}}$. It will be covariant if these are chosen to transform like scalar fields.

2 Here I'm using the notion of the infinitesimal generator of a symmetry transformation. If you don't know this concept, take a quick look at Appendix G, or consult one of the many excellent introductions to Lie groups [1–4].

This equation for the momentum-space wave function of “a particle localized at the origin” is not the same as the one we are used to from non-relativistic quantum mechanics. However, if we are in the non-relativistic regime where $|p| \ll m$ then the wave function reduces to $1/m$ times the non-relativistic formula. When relativity is taken into account, the localized particle appears to be spread out over a distance of order its Compton wavelength, $1/m = \hbar/(mc)$.

Our formula for the emission/absorption amplitude is thus covariant, but it poses the following paradox: *it is non-zero when the separation between the emission and absorption points is space-like.* The causal order of two space-like separated points is not Lorentz-invariant (Problem 2.1), so this is a real problem.

The only known solution to this problem is to impose a new physical postulate: every emission source could equally well be an absorption source (and vice versa). We will see the mathematical formulation of this postulate in the next chapter. Given this postulate, we define a total source by $J(x) = J_{\mathrm{E}}(x) + J_{\mathrm{A}}(x)$ and write an amplitude

$$\begin{aligned} A_{\mathrm{AE}} &= \int \mathrm{d}^4x\, \mathrm{d}^4y\, J(x)J(y) \int \frac{\mathrm{d}^3p}{2\omega_p(2\pi)^3}[\theta(x^0 - y^0)\mathrm{e}^{-\mathrm{i}p(x-y)} \\ &\qquad\qquad + \theta(y^0 - x^0)\mathrm{e}^{\mathrm{i}p(x-y)}] \\ &\equiv \int \mathrm{d}^4x\, \mathrm{d}^4y\, J(x)J(y)D_{\mathrm{F}}(x-y), \end{aligned} \tag{1.4}$$

where $\theta(x^0)$ is the Heaviside step function which is 1 for positive x^0 and vanishes for $x^0 < 0$. From now on we will omit the 0 superscript in the argument of these functions. This formula is manifestly Lorentz-covariant when $x - y$ is time-like or null. When the separation is space-like, the momentum integrals multiplying the two different step functions are equal, and we can add them, again getting a Lorentz-invariant amplitude. It is also consistent with causality. In any Lorentz frame, the term with $\theta(x^0 - y^0)$ is interpreted as the amplitude for a positive-energy particle to propagate forward in time, being emitted at y and absorbed at x. The other term has a similar interpretation as emission at x and absorption at y. Different Lorentz observers will disagree about the causal order when $x - y$ is space-like, but they will all agree on the total amplitude for any distribution of sources.

Something interesting happens if we assume that the particle has a conserved Lorentz-invariant charge, like electric charge. In that case, one would have expected to be able to correlate the question of whether emission or absorption occurred to the amount of charge transferred between x and y. Such an absolute definition of emission versus absorption is not consistent with the postulate that saved us from a causality paradox. In order to avoid it we have to make another, quite remarkable, postulate: every charge-carrying particle has an anti-particle of exactly equal mass and opposite charge. If this is true we will not be able to use charge transfer to distinguish between emission of a particle and absorption of an anti-particle. One of the great triumphs of quantum field theory is that this prediction is experimentally verified. The equality of particle and anti-particle masses has been checked to one part in 10^{18} [5].

Now let’s consider a slightly more complicated process in which the particle scatters from some external potential before being absorbed. Suppose that the potential is

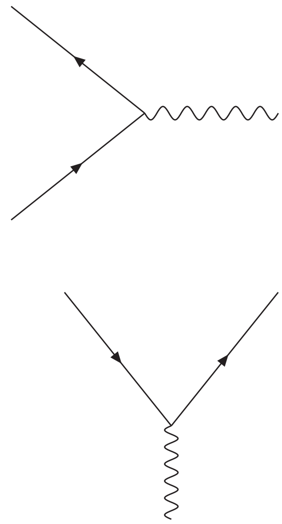

Fig. 1.2. Scattering in one frame is production amplitude in another.

short-ranged, and is turned on for only a brief period, so that we can think of it as being concentrated near a space-time point z. The scattering amplitude will be approximately given by propagation from the emission point to the interaction point z, some interaction amplitude, and then propagation from z to the absorption point. We can draw a space-time diagram like Figure 1.2. We have seen that the propagation amplitudes will be non-zero, even when all three points are at space-like separation from each other. Then, there will be some Lorentz frame in which the causal order is that given in the second drawing in the figure. An observer in this frame sees particles created from the vacuum by the external field! Scattering processes inevitably imply particle-production processes.[3]

We conclude that a theory consistent with special relativity, quantum mechanics, and causality must allow for particle creation when the energetics permits it (in the example of the previous paragraph, the time dependence of the external field supplies the energy necessary to create the particles). This, as we shall see, is equivalent to the statement that a causal, relativistic quantum mechanics must be a theory of quantized local fields. Particle production also gives us a deeper understanding of why the single-particle wave function is spread over a Compton wavelength. To localize a particle more precisely we would have to probe it with higher momenta. Using the relativistic energy–momentum relation, this means that we would be inserting energy larger than the particle mass. This will lead to uncontrollable pair production, rather than localization of a single particle.

Before leaving this introductory section, we can squeeze one more drop of juice from our simple considerations. This has to do with how to interpret the propagation amplitude $D_{\mathrm{F}}(x-y)$ when $x-y$ is space-like, and we are in a Lorentz frame where

[3] Indeed, there are quantitative relations, called *crossing symmetries*, between the two kinds of amplitude.

$x^0 = y^0$. Our remarks about particle creation suggest that we should interpret this as the probability amplitude for two particles to be found at time x^0, at relative separation $x - y$. Note that this amplitude is completely symmetric under interchange of x and y, which suggests that the particles are bosons. We thus conclude that spin-zero particles must be bosons. It turns out that this is true, and is a special case of a theorem that says that integer-spin particles are bosons and half-integer spin particles are fermions. What is more, this is not just a mathematical theorem, but an experimental fact about the real world. I should warn you, though, that unlike the other remarks in this section, and despite the fact that it leads to a correct conclusion, the reasoning here is not a cartoon of a rigorous mathematical argument. The interpretation of the equal-time propagator as a two-particle amplitude is of limited utility.

2 Quantum theory of free scalar fields

V. Fock invented an efficient method for dealing with multiparticle states. We will work with delta-function-normalized single-particle states in describing Fock space. This has the advantage that we never have to discuss states with more than one particle in exactly the same state, and various factors involving the number of particles drop out of the formulae. In this section we will continue to work with spinless particles.

Start by defining Fock space as the direct sum $\mathcal{F} = \bigoplus_{k=0}^{\infty} \mathcal{H}_k$, where $\mathcal{H}_k$ is the Hilbert space of k particle states. We will assume that our particles are either bosons or fermions, so these states are either totally symmetric or totally anti-symmetric under particle interchange. In particular, if we work in terms of single-particle momentum eigenstates, $\mathcal{H}_k$ consists of states of the form $|p_1, \ldots, p_k\rangle$, either symmetric or anti-symmetric under permutations. The inner product of two such states is

$$\langle p_1, \ldots, p_k | q_1, \ldots, q_l \rangle = \delta_{kl} \frac{1}{k!} \sum_{\sigma} (-1)^{S\sigma} \delta^3(p_1 - q_{\sigma(1)}) \ldots \delta^3(p_k - q_{\sigma(k)}), \tag{2.1}$$

where the sum is over all permutations σ in the symmetric group, S_k, $(-1)^{\sigma}$ is the sign of the permutation, and the statistics factor, S, is 0 for bosons and 1 for fermions.

The $k = 0$ term in this direct sum is a one-dimensional Hilbert space containing a unique normalized state, called the vacuum state and denoted by $|0\rangle$.

In ordinary quantum mechanics one can contemplate particles that form different representations of the permutation group than bosons or fermions. Although we will not prove it in general, this is impossible in quantum field theory. In one or two spatial dimensions, one can have particles with different statistical properties (braid statistics), but these can always be thought of as bosons or fermions with a particular long-range interaction. In three or more spatial dimensions, only Bose and Fermi statistics are allowed for particles in Lorentz-invariant QFT.

Fock realized that one can organize all the multiparticle states together in a way that simplifies all calculations. Starting with the normalized state $|0\rangle$, which has no particles in it, we introduce a set of commuting, or anti-commuting, operators $a^{\dagger}(p)$ and define

$$|p_1, \ldots, p_k\rangle \equiv a^{\dagger}(p_1) \ldots a^{\dagger}(p_k)|0\rangle. \tag{2.2}$$

These are called the creation operators. The scalar-product formula is reproduced correctly if we postulate the following (anti-)commutation relation between creation operators and their adjoints[1] (called the annihilation operators):

$$[a(p), a^{\dagger}(q)]_{\pm} = \delta^3(p - q). \tag{2.3}$$

Fermions are made with anti-commutators, and bosons with commutators.

To get a little practice with Fock space, let's construct the representation of the Poincaré symmetry[2] on the multiparticle Hilbert space. We begin with the energy and momentum. These are diagonal on the single-particle states. The correct Fock-space formula for them is

$$P^{\mu} = \int \mathrm{d}^3 p \, p^{\mu} a^{\dagger}(p) a(p), \tag{2.4}$$

where $p^0 = \omega_p$. Its easy to verify that this operator does indeed give us the sum of the k single-particle energies and momenta, when acting on a k-particle state. This is because $n_p \equiv a^{\dagger}(p)a(p)$ acts as the particle number density in momentum space. There is a similar formula for all operators that act on a single particle at a time. For example, the rotation generators are

$$J_{ij} = \int \mathrm{d}^3 p \, a^{\dagger}(p) i(p_i \, \partial_j - p_j \, \partial_i) a(p). \tag{2.5}$$

Here $\partial_j = \partial / \partial p^j$.

It is easy to verify that the following formula defines a unitary representation of the Lorentz group on single-particle states:

$$U(\Lambda)|p\rangle = \sqrt{\frac{\omega_{\Lambda p}}{\omega_p}} |\Lambda p\rangle.$$

The reason for the funny factor in this formula is that the Dirac delta function in our definition of the normalization is not covariant because it obeys

$$\int \mathrm{d}^3 p \, \delta(p) = 1.$$

A Lorentz-invariant measure of integration on positive-energy time-like 4-vectors is

$$\int \mathrm{d}^4 p \, \theta(p_0) \delta(p^2 - m^2) = \int \frac{\mathrm{d}^3 p}{2\omega_p}.$$

The factors in the definition of $U(\Lambda)$ make up for this non-covariant choice of normalization.

A general Lorentz transformation is the product of a rotation and a boost, so in order to complete our discussion of Lorentz generators we have to write a formula for

[1] This is an extremely important claim. It's easy to prove and every reader should do it. The same remark applies to all of the equations in this subsection. Commutators and anti-commutators of operators are defined by $[A, B]_{\pm} \equiv AB \pm BA$. Assume that $a(p)|0\rangle = 0$.

[2] Poincaré symmetry is the semi-direct product of Lorentz transformations and translations. Semi-direct means that the Lorentz transformations act on the translations.

the generator of a boost with infinitesimal velocity v. We write it as $v^i J_{0i}$. Under such a boost, $p^i \to p^i + v^i \omega_p$ and $\omega_p \to \omega_p + v^i p_i$. Thus

$$\delta|p\rangle = \left(\frac{p_i v^i}{2\omega_p} + \omega_p v^i \frac{\partial}{\partial p^i}\right)|p\rangle.$$

The Fock-space formula for the boost generator is then

$$v^i J_{0i} = \int \mathrm{d}^3 p \left[a^\dagger(p)\left(\frac{p_i v^i}{2\omega_p} + \omega_p v^i \frac{\partial}{\partial p^i}\right)a(p)\right].$$

2.1 Local fields

We now want to model the response of our infinite collection of scalar particles to a localized source $J(x)$. We do this by adding a term to the Hamiltonian (in the Schrödinger picture)

$$H \to H_0 + V(t), \tag{2.6}$$

with

$$V(t) = \int \mathrm{d}^3 x\, \phi(x) J(x, t). \tag{2.7}$$

$\phi(x)$ must be built from creation and annihilation operators. It must transform into $\phi(x+a)$ under spatial translations. This is guaranteed by writing $\phi(x) = \int \mathrm{d}^3 p\, e^{ipx} \hat{\phi}(p)$, where $\hat{\phi}(p)$ is an operator carrying momentum p.

We want to model a source that creates and annihilates single particles. This statement is meant in the sense of perturbation theory. That is, the amplitude J for the source to create a single particle is small. It can create multiple particles by multiple action of the source, which will be higher-order terms in a power series in J. Thus the field $\phi(x)$ should be linear in creation and annihilation operators:

$$\phi(x) = \int \frac{\mathrm{d}^3 p}{\sqrt{(2\pi)^3 2\omega_p}} \left[a(p)\alpha \mathrm{e}^{\mathbf{ipx}} + a^\dagger(p)\alpha^* \mathrm{e}^{-\mathbf{ipx}}\right]. \tag{2.8}$$

We have also imposed Hermiticity of the Hamiltonian, assuming that the source function is real.[3] α could be a general complex constant, but we have already defined a normalization for the field, and we can absorb the phase of α into the creation and annihilation operators, so we set $\alpha = 1$.

We now want to study the time development of our system in the presence of the source. We are not really interested in the free motion of the particles, but rather in the question of how the source causes transitions between eigenstates of the free Hamiltonian. Dirac invented a formalism, called the Dirac picture, for studying problems of

[3] As we will see, complex sources are appropriate for the more complicated situation of particles that are not their own anti-particles.