How to integrate out a particle

Vanessa Alexandra Hollmeier

September 2024

TL;DR: I show how to integrate out particles in the QFT path integral formalism analytically, and visually explain what happens in terms of Feynman diagrams.

Contents

1 Introduction

Ever so often, we hear the phrase "integrating out a particle" in QFT. Roughly speaking, it refers to a procedure in which we remove a particle from our theory, and simultaneously modify the action of our theory such that the rest of the theory still looks the same. However, I've never seen an actual, high-level systematic explanation on how to do this anywhere, so I'm writing this blog post now.

1.1 Basics

We are going to use a simple toy model with a phion and chion real scalar field:

$$S[\hat{\phi}, \hat{\chi}] = S_\phi[\hat{\phi}] + S_\chi[\hat{\chi}] + S_{\text{int}}[\hat{\phi}, \hat{\chi}], \tag{1}$$

where

$$S_\phi[\hat{\phi}] = \frac{1}{2}(\partial_\mu \hat{\phi})(\partial^\mu \hat{\phi}) - \frac{1}{2} m_\phi^2 \hat{\phi}^2 \tag{2}$$

$$S_\chi[\hat{\chi}] = \frac{1}{2}(\partial_\mu \hat{\chi})(\partial^\mu \hat{\chi}) - \frac{1}{2} m_\chi^2 \hat{\chi}^2, \tag{3}$$

and $S_{\text{int}}[\hat{\phi}, \hat{\chi}]$ is some interaction vertex between them.

The hat notation I'm using here isn't supposed to mean that the fields are operators, but rather that they are fields that haven't been integrated over in the path integral yet - they are, so to speak, the path integral analogue of operators.

In path integral QFT, we now define the path integral

$$\int \mathcal{D}\hat{\phi}\mathcal{D}\hat{\chi} \exp\Big(iS[\hat{\phi}, \hat{\chi}]\Big) \tag{4}$$

If we want to extract predictions from it, we add source currents J, K and define the generating functionals[1]

$$Z[J, K] = \int \mathcal{D}\hat{\phi}\mathcal{D}\hat{\chi} \exp\Big(i(S[\hat{\phi}, \hat{\chi}] + J\hat{\phi} + K\hat{\chi})\Big) \tag{5}$$

$$W[J, K] = i \ln Z[J, K] \tag{6}$$

such that when we take a functional derivative

$$\frac{\delta^{n+k} W[J, K]}{\delta^n J \; \delta^k K} \tag{7}$$

we obtain the sum of all Feynman diagrams with n external phion legs and k external chion legs. Basically, in intuitive terms, we wiggle around the source currents to check how they react to the fields.

Now, suppose that we are **not interested in chions at all** - for instance because our accelerator only accelerates phions, and our detectors only detect phions. This means that we're not interested in Feynman diagrams with external chion legs - the only chions appearing in our diagrams are internal propagators. Basically, we fix the chion source current K to zero everywhere and vow to never touch it again:

$$Z'[J] = Z[J, K = 0] \tag{8}$$

$$W'[J] = W[J, K = 0] \tag{9}$$

As before, we can take derivatives of $W'[J]$ and get all Feynman diagrams with phion external legs.

If we want to calculate these Feynman diagrams, we now have two options:

[1] I am using the Einstein integration convention, i.e. $J\phi = \int \mathrm{d}^4 x J(x)\phi(x)$.

- We proceed exactly as before - we take all the internal chion propagators into account and sum them up.
- We try to find a new action $S'[\hat{\phi}]$ with new parameters, such that:

$$Z'[J] =: \int \mathcal{D}\hat{\phi} \exp\Big(i(S'[\hat{\phi}] + J\hat{\phi})\Big). \tag{10}$$

In other words - we're trying to find a new theory $S'[\hat{\phi}]$ that will yield **the exact same predictions for phion interactions** as our previous interacting theory - except that it **doesn't contain chions**. The effects that the internal chion propagators had on pure phion interactions should be accounted for by shifting around the parameters of the old phions and introducing new vertices. When we have done so, we say that we have **integrated out the chion**.

Normally, we do this by performing only the chion part of the partial integral. Let's look at the simplest possible example first:

1.2 The easiest case

The easiest case is a theory in which the phions and chions don't interact at all, i.e. $S_{\text{int}}[\hat{\phi}, \hat{\chi}] = 0$:

$$S[\hat{\phi}, \hat{\chi}] = S_\phi[\hat{\phi}] + S_\chi[\hat{\chi}] \tag{11}$$

Visually speaking: The phions do their thing, and the chions do their thing, but they don't influence each other - and when we integrate out the chions, we expect that the interaction between the phions stay exactly the same.

In fact, the path integral nicely decomposes into a phion part and a chion part:

$$Z[J, K] = \int \mathcal{D}\hat{\phi}\mathcal{D}\hat{\chi} \exp\Big(i(S_\phi[\hat{\phi}] + S_\chi[\hat{\chi}] + J\hat{\phi} + K\hat{\chi})\Big) \tag{12}$$

$$= \int \mathcal{D}\hat{\phi}\mathcal{D}\hat{\chi} \exp\Big(i(S_\phi[\hat{\phi}] + J\hat{\phi})\Big) \exp(i(S_\chi[\hat{\chi}] + K\hat{\chi})) \tag{13}$$

$$= \int \mathcal{D}\hat{\phi} \exp\Big(i(S_\phi[\hat{\phi}] + J\hat{\phi})\Big) \int \mathcal{D}\hat{\chi} \exp(i(S_\chi[\hat{\chi}] + K\hat{\chi})) \tag{14}$$

$$= \left(\int \mathcal{D}\hat{\phi} \exp\Big(i(S_\phi[\hat{\phi}] + J\hat{\phi})\Big)\right) \left(\int \mathcal{D}\hat{\chi} \exp(i(S_\chi[\hat{\chi}] + K\hat{\chi}))\right) \tag{15}$$

$$\tag{16}$$

Now, when we fix $K = 0$, the second term is just some constant $\mathcal{N}$, and we have found our $S'[\hat{\phi}]$:

$$Z'[J] = Z[J, K = 0] = \mathcal{N} \left(\int \mathcal{D}\hat{\phi} \exp\Big(i(S_\phi[\hat{\phi}] + J\hat{\phi})\Big)\right) \tag{17}$$

$$\implies S'[\hat{\phi}] = S_\phi[\hat{\phi}] \tag{18}$$

However, when the two particles interact, the new action S' is bound to be something different. In the following, I will explain the two cases where is is possible to integrate out the chion in an interacting theory analytically.

2 Integrating out linear interactions

First of all, we consider the case where the interaction vertex $S_{\text{int}}[\hat{\phi}, \hat{\chi}]$ is linear in χ - or in other words, the interaction vertex has exactly one χ arm. For instance:

$$S_{\text{int}}[\hat{\phi}, \hat{\chi}] = -g\hat{\chi}\hat{\phi}^2 \tag{19}$$

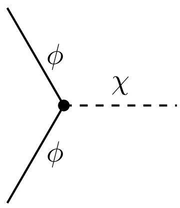

Now, let's assume that we only have phion external legs. The only way that internal chion lines can influence the Feynman diagram is by linking two of the vertices together. In the integrated-out theory, we'd therefore expect a contracted 4-phion vertex:

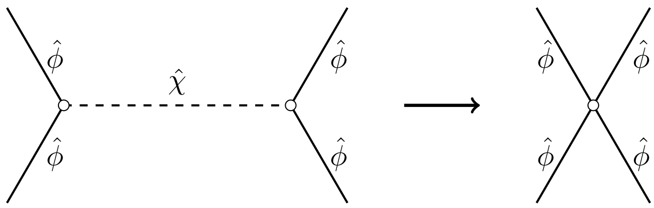

How does this play out mathematically? First, we split up the action of our theory like this:

$$S[\hat{\phi}, \hat{\chi}] = S_\phi[\hat{\phi}] + \frac{1}{2}\hat{\chi}(-\partial^2 - m_\chi^2)\hat{\chi} - g\hat{\chi}\hat{\phi}^2 \tag{20}$$

For simplicity, we write the kinetic operator as $A = (-\partial^2 - m^2)$. We now complete the square:

$$S[\hat{\phi}, \hat{\chi}] = S_\phi[\hat{\phi}] + \frac{1}{2}\hat{\chi}A\hat{\chi} - g\hat{\chi}\hat{\phi}^2 \tag{21}$$

$$= S_\phi[\hat{\phi}] + \frac{1}{2}\hat{\chi}A\hat{\chi} \tag{22}$$

$$- \frac{1}{2}g\hat{\chi}AA^{-1}\hat{\phi}^2 \tag{23}$$

$$- \frac{1}{2}g\hat{\phi}^2A^{-1}A\hat{\chi} \tag{24}$$

$$+ \frac{1}{2}g^2\hat{\phi}^2A^{-1}AA^{-1}\hat{\phi}^2 \tag{25}$$

$$- \frac{1}{2}g^2\hat{\phi}^2A^{-1}AA^{-1}\hat{\phi}^2 \tag{26}$$

$$= S_\phi[\hat{\phi}] + \frac{1}{2}\left(\hat{\chi} - g\hat{\phi}^2A^{-1}\right)A\left(\hat{\chi} - gA^{-1}\hat{\phi}^2\right) - \frac{1}{2}g^2\hat{\phi}^2A^{-1}\hat{\phi}^2 \tag{27}$$

Now, we can insert this action into the path integral:

$$\int \mathcal{D}\hat{\phi}\mathcal{D}\hat{\chi}\exp\Bigg(i\left(S_\phi[\hat{\phi}] + J\hat{\phi} + \left(\hat{\chi} - g\hat{\phi}^2A^{-1}\right)A\left(\hat{\chi} - gA^{-1}\hat{\phi}^2\right) - \frac{1}{2}g^2\phi^2A^{-1}\phi^2\right)\Bigg) \tag{28}$$

$$= \int \mathcal{D}\hat{\phi}\exp\Big(iS_\phi[\hat{\phi}] + J\hat{\phi}\Big)\left(\int \mathcal{D}\hat{\chi}\exp\Bigg(i\left(\left(\hat{\chi} - g\hat{\phi}^2A^{-1}\right)A\left(\hat{\chi} - gA^{-1}\hat{\phi}^2\right) - \frac{1}{2}g^2\phi^2A^{-1}\phi^2\right)\Bigg)\right) \tag{29}$$

We shift the χ field integration variable in the inner integral to make it independent of ϕ, i.e. $\chi \to \chi - gA^{-1}\phi^2$. Then, the inner integral suddenly gets a lot simpler:

$$Z[J, K = 0] = \int \mathcal{D}\hat{\phi}\exp\Big(iS_\phi[\hat{\phi}] + J\hat{\phi}\Big)\left(\int \mathcal{D}\hat{\chi}\exp\Bigg(i\left(\hat{\chi}A\hat{\chi} - \frac{1}{2}g^2\phi^2A^{-1}\phi^2\right)\Bigg)\right) \tag{30}$$

The second part of the second exponent is entirely independent of χ, so we can pull it out of the second integral:

$$Z[J, K = 0] = \int \mathcal{D}\hat{\phi}\exp\Bigg(iS_\phi[\hat{\phi}] - \frac{1}{2}g^2\phi^2A^{-1}\phi^2 + J\hat{\phi}\Bigg)\left(\int \mathcal{D}\hat{\chi}\exp(i\left(\hat{\chi}A\hat{\chi}\right))\right) \tag{31}$$

In turn, the remaining part of the second integral is independent of ϕ, so it's just an irrelevant constant factor $\mathcal{N}$ in front of the path integral:

$$Z[J, K = 0] = \mathcal{N}\int \mathcal{D}\hat{\phi}\exp\Bigg(iS_\phi[\hat{\phi}] - \frac{1}{2}g^2\phi^2A^{-1}\phi^2 + J\hat{\phi}\Bigg) \tag{32}$$

Now, we can directly read off the new integrated-out action:

$$S'[\hat{\phi}] = S_\phi[\hat{\phi}] - \frac{1}{2}g^2\phi^2\frac{1}{-\partial^2 - m_\chi^2}\phi^2. \tag{33}$$

That's exactly what we expected! Or is it? We should ask ourselves two very serious questions right now:

- Before we integrated out the chion, the chion-mediated interaction was **nonlocal**: Figuratively speaking, two phions collide to produce a chion, the chion moves somewhere else, and then decays into two new phions. But a literal four-phion vertex should be **local**: Two phions collide and form two new phions at the same spot. So how does that fit together? The integrated-out theory is supposed to give the same predictions for phions!
- How the heck are we supposed to interpret the propagator sticking out in the middle of the four-phion vertex above?

You might've already guessed it - the two questions answer each other. To see why, we need to descend from our notational ivory tower for a second and make the involved coordinates explicit.

First of all, we are going to denote the chion propagator by:

$$G(y, x) = \frac{1}{-\partial^2 - m_\chi^2} \tag{34}$$

The semantic content of this object is: "If we create a chion particle at time-position x, we are going to find it at time-position y with the amplitude $G(y, x)$". The chion *propagator* describes how the chion particle *propagates* - and in turn, it also describes the non-locality of the chion-mediated phion interaction.

The four-phion vertex from (33) has the following written-out form:

$$-\frac{1}{2}g^2\int \mathrm{d}^4y\mathrm{d}^4x\ \hat{\phi}^2(y)\ G(y, x)\ \hat{\phi}^2(x) \tag{35}$$

There's no chion field in this vertex - but we can see how the vertex "emulates" the effects of the chion without actually including it: Two phions collide at time-position x, and two new pions are created at time-position y - with the amplitude with which a chion *would* propagate from x to y. Pretty cool, isn't it?

3 Integrating out quadratic interactions

The story gets a lot more interesting once we allow interactions that are *quadratic* in χ, i.e. vertices with exactly two chion arms. For example:

$$S_{\text{int}}[\hat{\phi}, \hat{\chi}] = -g\hat{\chi}^2\hat{\phi} \tag{36}$$

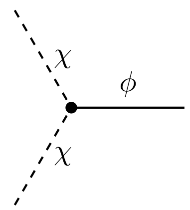

Again, let's assume that we only care about Feynman diagrams with phion external legs. As before, the only way chions can influence these diagrams is by internal propagators. So let's think about what shapes these propagators can take.

Let's say that there is an internal chion line somewhere in the Feynman diagram. It has two ends - but when they connect to another propagator, the line does not end. Instead, it leaves the vertex again, and goes on until another vertex, and so on. Because we assumed that our diagram only has external phion lines, it follows that the internal chion chain must eventually bite its own tail and become a loop. Therefore, all internal chions in these diagrams are chion loops.

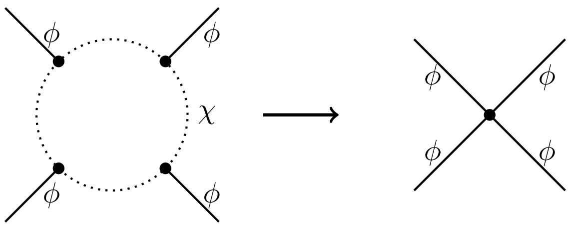

Again, when we integrate out the chion, we'll have to find a way to describe the effects of all possible chion loops on the phions. We will need to add one vertex with k phions for every chion loop with k phion arms.

At first glance, this looks like an impossible problem - but upon closer inspection, it tells us a lot about (a) the meaning of functional determinants, (b) the difference between bosonic and fermionic loops, and (c) the nature of gauge ghosts.

We start by writing out the action of the original theory the following way:

$$S[\hat{\phi}, \hat{\chi}] = S_\phi[\hat{\phi}] + \frac{1}{2}\hat{\chi}(-\partial^2 - m_\chi^2)\hat{\chi} - g\hat{\chi}^2\hat{\phi} \tag{37}$$

This time, we combine the $\chi\chi\phi$ vertex with the chion kinetic term:

$$S[\hat{\phi}, \hat{\chi}] = S_\phi[\hat{\phi}] + \frac{1}{2}\hat{\chi}(-\partial^2 - m_\chi^2 - g\hat{\phi})\hat{\chi} \tag{38}$$

The path integral now reads

$$Z[J, K = 0] = \int \mathcal{D}\hat{\phi}\mathcal{D}\hat{\chi} \exp\Big(i(S_\phi[\hat{\phi}] + J\hat{\phi} + \frac{1}{2}\hat{\chi}(-\partial^2 - m_\chi^2 - g\hat{\phi})\hat{\chi})\Big) \tag{39}$$

The first part is independent of χ, so we can pull it out of the chion path integral:

$$Z[J, K = 0] = \int \mathcal{D}\hat{\phi} \exp\Big(i(S_\phi[\hat{\phi}] + J\hat{\phi})\Big) \int \mathcal{D}\hat{\chi} \exp\Big(i\frac{1}{2}\hat{\chi}(-\partial^2 - m_\chi^2 - g\hat{\phi})\hat{\chi}\Big) \tag{40}$$

This time however, the second part is *not* independent of ϕ, so we can't just split the path integral into two. We actually have to pay attention to the result of the chion path integral now, instead of dismissing it as an unimportant constant!

If you remember your path integral QFT classes, you'll know that:

$$\int \mathcal{D}\hat{\chi} \exp(i\chi A\chi) = \mathcal{N}\sqrt{\frac{1}{\det A}} \tag{41}$$

where $\mathcal{N}$ is some constant. It follows that:

$$Z[J, K = 0] = \mathcal{N} \int \mathcal{D}\phi \exp\Big(i(S_\phi[\hat{\phi}] + J\hat{\phi})\Big) \Big(\det\Big(-\partial^2 - m_\chi^2 - g\hat{\phi}\Big)\Big)^{-1/2} \tag{42}$$

We can pull the determinant up into the exponential by taking its logarithm:

$$Z[J, K = 0] = \mathcal{N} \int \mathcal{D}\phi \exp\Big(i\Big(S_\phi[\hat{\phi}] + J\hat{\phi} + \frac{i}{2}\ln\det\Big(-\partial^2 - m_\chi^2 - g\hat{\phi}\Big)\Big)\Big) \tag{43}$$

So now, we can see that the integrated-out action is:

$$S'[\hat{\phi}] = S_\phi[\hat{\phi}] + \frac{i}{2}\ln\det\Big(-\partial^2 - m_\chi^2 - g\hat{\phi}\Big) \tag{44}$$

Okay. Now that's definitely NOT what we're used to seeing in a classical action. Normally, the classical actions of our theories just consist of polynomials of the fields. We know how to translate them to Feynman diagrams. But the logarithm of a determinant of a matrix? What does that even mean? How do we build Feynman diagrams from that? To find out, we'll have to perturbatively expand that term.

Let's treat the $-g\hat{\phi}$ bit as a pertubation around the free chion propagator:

$$A = -\partial^2 - m_\chi^2 \tag{45}$$

$$\delta A = -g\hat{\phi} \tag{46}$$

Now, we'd like to perturbatively expand $\ln\det(A+\delta A)$. For that, we need to calculate the first functional derivative of $\ln\det A$ - we vary A to $A+\delta A$ and look at how $\delta \ln\det A$ depends on δA:

$$\delta \ln\det A = \delta \,\mathrm{tr}\ln A = \mathrm{tr}\,\delta \ln A = \mathrm{tr}\left(\frac{1}{A}\delta A\right) \tag{47}$$

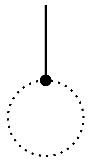

If we translate this to Feynman diagram notation, we see that this vertex of the new theory looks exactly like a chion loop term with one external phion leg. Taking another derivative will act on the propagator $1/A$ inside the trace:

$$\delta\frac{1}{A} = -\frac{1}{A}\delta A\frac{1}{A} \tag{48}$$

so for the second functional derivative of $\ln\det A$, we get:

$$\delta^2 \ln\det(A) = -\,\mathrm{tr}\ln\left(\frac{1}{A}\delta A\frac{1}{A}\delta A\right) \tag{49}$$

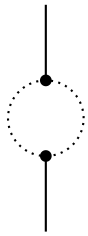

- the second loop term! And so on. By taking all possible functional derivatives of the $\ln\det$ term, we get all possible loop terms **as vertices**. The last part is crucial - our new theory **doesn't actually contain chions**, the perturbatively expanded vertices **just look exactly like the chion loops**:

$$S[\hat{\phi},\hat{\chi}] = S_\phi[\hat{\phi}] - \mathcal{N} + g\frac{i}{2}\,\mathrm{tr}\left(\frac{1}{-\partial^2 - m_\chi^2}\hat{\phi}\right) - g^2\frac{i}{2}\,\mathrm{tr}\left(\frac{1}{-\partial^2 - m_\chi^2}\hat{\phi}\frac{1}{-\partial^2 - m_\chi^2}\hat{\phi}\right) + \cdots \tag{50}$$

This way, it is guaranteed that the integrated-out theory for phions predicts the same things as the full coupled phion-chion theory.

As previously mentioned, I think that this technique can teach us a lot beyond the question of integrating out stuff:

(a) The $\ln\det A = \mathrm{tr}\ln A$ term generates loop terms. This is very useful to keep in mind while deriving the one-loop effective potential, aka the Coleman-Weinberg potential.

(b) If we integrate out a bosonic kinetic term, we get a $\sqrt{1/\det A}$ term. If we integrate out a fermionic kinetic term, however, we get a $\det A$ term. If we take the logarithm of these two, we can see that they have opposite signs. This sign difference is responsible for the fact that fermionic loops have an extra minus sign in front of them[2].

(c) In this post, we've turned the path integral over a field into a functional determinant that generates k-phion vertices which look exactly like chion loop terms. In the Fadeev-Popov procedure for gauge-fixing non-abelian gauge theories, we do the **exact opposite** - there, the determinant of the gauge fixing operator looks exactly like the result of the integrating-out of a fermion, so we invent a new fermion called the "gauge ghost", which, when integrated out, yields the exact determinant/loop vertices we had in the beginning. In other words, we "integrate in" the gauge ghost. Of course, hell breaks loose when you then add a source current for this "particle" and then wonder why you have unphysical in and out states[3]. Who could've guessed. So, the next time you think about gauge ghosts, remember that they are not real - the only thing that's real here are the multi-gluon vertices that just happen to look like they're fermionic loop terms of some fictious gauge ghost.

4 Conclusion

- Integrating out particles with vertices linear in them results in new contracted vertices containing a propagator of the integrated-out particle.
- Integrating out particles with vertices quadratic in them results in new vertices that look like loop terms of the integrated-out particle.

[2] which, in turn, justifies the Grassmann number integration rules for fermions, I think.
[3] But somehow, everyone seems to do that anyway...? If you know why, hit me up.