Introduction to finite-temperature field theory

Saga Säppi

July 19, 2023

1 Motivation

The main question finite-temperature field theory (FTFT) [1] tries to answer is the following:

How do quantum fields behave in a thermal setting?

To attempt to answer this question, we are going to make an important assumption: That the fields in question exist in a *thermal bath* characterised by a temperature T. This means that they exist in a background of other particles or fields that have already equilibriated in an unspecified way so that a well-defined temperature exists. Therefore, at most, we can consider small deviations of this equilibrium and see what their effects are. As an upside, many approach of classical thermodynamics and statistical mechanics that you have learned in undergraduate studies will translate to FTFT with little effort. It is possible to consider also out-of-equilibrium systems where a well-defined temperature does not exist, but this is generally much more difficult and there is no well-established and complete way of doing it; it is very much done on a case-by-case basis.

Now, what kind of thermal quantities one can calculate in equilibriated systems? Preferably, these should also be something that can be measured.

In true equilibrium, such quantities are found in thermodynamics. The most basic one is something you've probably seen in school: the free energy density (minus the pressure), or, in statistical parlance, $\Omega = -\ln Z$ where Z is the partition function: If you remember your statistical physics, you know that this is the fundamental quantity that codifies the information about a thermal system in equilbirium. Whether we are talking about a system of point-like particles or a quantum field as the medium in which the pressure is considered, this quantity will have a very similar interpretation. With the partition function known, we can obtain other quantities you might remember from Thermo or StatMech: The energy density $\varepsilon = T^2 \partial_T \ln Z$, the entropy $s = \ln Z + \varepsilon/T$, or the speed of sound in the medium characterised by the field $c_s^2 = s/\partial_T \varepsilon = -\partial_\varepsilon \Omega$. At finite chemical potentials, you would also have eg. susceptibilities $\chi_{ij} = \partial_{\mu_i} \partial_{\mu_j} \Omega$. A common factor here is that all these quantities are time-independent. This will become crucial soon, when we develop the so-called *imaginary-time formalism* for computing them in field theory.

However, there are also time-dependent quantities of interest. These are associated with small deviations from the equilibrium, in such a way that a well-defined temperature still exists. Examples of such quantities are correlation functions, which in a quantum field theory are correlation functions of fields—say, $\langle \phi(x) \phi(y) \rangle$ where x, y are space*time* points—and as such, some of the most fundamental objects in the theory. More tangible quantities include, for example, the rates at which particles are produced in a medium, various damping rates which characterise how the medium affects the propagation of field modes through it, and transport quantities, like viscosity. This will be the topic of the second main part of the lectures, in the context of the *real-time formalism*. If all goes well, I will end the lectures with a discussion on effective theories in a thermal context, in the spirit of the school.

Due to my own personal bias, these lecture notes will focus primarily on relativistic field theories, with quantum chromodynamics (QCD) as the main motivation. In practice, this means that the tempreatures are quite large (for QCD, we are talking at the very least about hundreds of MeVs, or trillions of Kelvins when T is the dominant scale[2]), and for the same reason I will also be using natural units ($\hbar = k_B = c = 1$) everywhere. Another curiosity regarding notation is the dimension: The dimension of the spacetime will be denoted with D, while the dimensionality of just the space-component is d. Usually, these are (close to) four and three respectively.

[1] Also Thermal Field Theory (TFT), although the acronym should be used with caution since it coincides with Topological Field Theory (also TFT).

[2] And even in "cold dense systems" like neutron stars discussed before, we are typically in billions of Kelvins, which is 'basically nothing' :)

With all that said, thermal field theory has a plethora of applications elsewhere; in condensed matter, nuclear physics, and such, where one often talks about *statistical field theory*. Almost identical methods are also useful in systems with other external scales: We will see (and you might have already seen) how a chemical potential can be turned on and included as a thermal scale analogously to classical thermodynamics, but for example strong external (electric or magnetic) fields also have a very similar effect to those of a temperature! Another glaring omission made in these notes is the absence of any discussion on Wilson loops, and consequently phase transitions or the QCD phase diagram; I hope that the lattice-FTFT lectures will cover the topic, as it is a fundamentally nonperturbative one.

Literature

- Kapusta & Gale
- Le Bellac
- Vuorinen & Laine
- Zinn-Justin
- GKSV

2 Imaginary-time formalism

2.1 Defining and motivating the formalism

In standard statistical mechanics, the partition function can be computed as $Z = \mathrm{Tr}e^{-\beta H}$, where Tr is some appropriate definition of a trace, $\beta = 1/T$, and H a (time-independent) Hamiltonian. Correlation values are similarly obtained with

$$\langle \mathcal{O} \rangle_{\text{statmech}} = Z^{-1}\mathrm{Tr}\left[\mathcal{O}e^{-\beta H}\right],\ Z = \mathrm{Tr}\left[e^{-\beta H}\right]$$

where $\mathcal{O}$ is some observable (say, a Hermitian operator in quantum statistical mechanics). This bears a striking resemblance to the path-integral formalism of QFT of fields Φ with a Euclidean Lagrangian density $\mathcal{L}_E$. Then, the correlation function of an observable $\mathcal{O}$ is obtained as the Euclidean path integral

$$\langle \mathcal{O} \rangle_{T=0\text{ QFT}} = Z^{-1}\int \mathscr{D}\Phi\mathcal{O}\exp\left[-\int_x \mathrm{d}^D x\mathcal{L}_E\left[\Phi\right]\right] \equiv Z^{-1}\mathrm{Tr}\left[\mathcal{O}e^{-S_E}\right]\ Z = \int \mathscr{D}\Phi\exp\left[-\int_x \mathrm{d}^D x\mathcal{L}_E\left[\Phi\right]\right].$$

For a thermal QFT, I make the following claim:

Given a time-independent observable $\mathcal{O}$ in a thermal QFT characterised by a set of fields Φ, its correlator function can be obtained as the Euclidean path integral

$$\langle \mathcal{O} \rangle_{\text{finite}-T\text{ QFT}} \equiv \langle \mathcal{O} \rangle = Z^{-1}\int \mathscr{D}\Phi\mathcal{O}\exp\left[-\int_0^\beta \mathrm{d}t\int_{\boldsymbol{x}} \mathrm{d}^d\boldsymbol{x}\mathcal{L}_E\left[\Phi\right]\right],\ Z = \int \mathscr{D}\Phi\exp\left[-\int_0^\beta \mathrm{d}t\int_{\boldsymbol{x}} \mathrm{d}^d\boldsymbol{x}\mathcal{L}_E\left[\Phi\right]\right]$$

That is to say, a finite-temperature field theory is simply a vacuum field theory defined on a base space whose time-factor has undergone a Kaluza–Klein compactification into S^1_β, a circle of length β. [3]

[3] If this phrasing makes more sense to you, I recommend looking up things on ncatlab or, for a gentler start, reading Baez & Muniain, even though it does not discuss thermal field theory specifically.

This section is dedicated to making somewhat hand-wavy arguments for this claim (as well as looking into some subtleties, like boundary conditions) and applying it to some simple examples using perturbation theory. Perhaps the simplest method is to note that, at least for a canonical Hamiltonian, it holds that the amplitude between two states separated by a time t

$$\left\langle \Phi_f \left| e^{-iHt} \right| \Phi_i \right\rangle_{T=0\text{ QFT}} = \int_{\Phi(0)=\Phi_i}^{\Phi(t)=\Phi_f} \mathscr{D}\Phi e^{iS},$$

where the integration space obeys appropriate boundary conditions. If we now assume the states i, f are evaluated at the same time t and analytically continute both sides to Euclidean space and, we obtain

$$\left\langle \Phi_f \left| e^{-\beta H} \right| \Phi_i \right\rangle_{T=0\text{ QFT}} = \int_{\Phi(t)=\Phi_i}^{\Phi(t-i\beta)=\pm\Phi_f} \mathscr{D}\Phi e^{-S_E}$$

where S_E is the Euclidean action for the configurations between $t = 0$ and $t = \beta$. The (anti)periodicity then arises by noting that the field configurations must agree up to a sign at the endpoints of the integration since we are computing the trace: We get $\mathrm{Tr}e^{-\beta H} = \int \mathrm{d}\Phi_i \int_{\Phi(t)=\Phi_i}^{\Phi(t-i\beta)=\pm\Phi_i} \mathscr{D}\Phi e^{-S_E}$, where the functional integral is compactly written as $\int \mathscr{D}\Phi e^{-S_E}$, indicating an integral over all field configurations which agree at the endpoints. The claim then follows immediately by noting that $\langle \mathbb{O} \rangle = \frac{1}{Z}\mathrm{Tr}\mathbb{O}e^{-\beta H}$ where the normalisation factor $Z = \mathrm{Tr}e^{-\beta H}$ is fixed by requiring $\langle \mathbb{O} \rangle = 1$. Another way to see the periodic boundaries arise is to consider a generic correlator of two operators $\{\mathbb{O}_i^M\}_{i\in\{1,2\}}$, which may or may not be composite, and which share a grading (both are either bosonic or fermionic), in *Minkowski* space. For simplicity, you could take these to be a set of scalar fields; we also suppress any possible external indices here as they do not play any role, and to start, we will let them depend on the time as well.

$$\left\langle \mathbb{O}_1^M(t_1, \boldsymbol{x}_1)\, \mathbb{O}_2^M(t_2, \boldsymbol{x}_2) \right\rangle = Z^{-1}\mathrm{Tr}\left[\mathbb{O}_1^M(t_1, \boldsymbol{x}_1)\, \mathbb{O}_2^M(t_2, \boldsymbol{x}_2)\, e^{-\beta H}\right]$$

$$= Z^{-1}\mathrm{Tr}\left[\mathbb{O}_1^M(t_1, \boldsymbol{x}_1)\, e^{-\beta H}\mathbb{O}_2^M(t_2 - i\beta)\right] = \left\langle \mathbb{O}_2^M(t_2 - i\beta, \boldsymbol{x}_1)\, \mathbb{O}_1^M(t_1, \boldsymbol{x}_2) \right\rangle.$$

For static observables, it then follows that the time-coordinate is either periodic or antiperiodic. This is in contrast to the vacuum case, where we at least formally typically allow all sufficiently quickly vanishing smooth field configurations. The relation derived here, especially in Minkowski space, goes by the name of a Kubo–Martin–Schwinger relation, and will be important for deriving the real-time formalism, as well.

2.2 Consequences for computations

Let us now consider some consequences of the formalism. An immediate consequence is that we have broken Lorentz symmetry: from $\mathrm{SO}(D)$, only a spatial $\mathrm{SO}(d)$ is left. This is natural, as we have singled out the time-direction in order to use it to encode temporal information. It also affects practical calculations: Spacetime tensor structures are no longer Lorentz-symmetric, only $\mathrm{SO}(d)$-symmetric, and there is now a new fundamental vector $N^\mu = (1, 0)$, which is the rest frame of the heat bath.

Another practical effect is seen when moving to momentum space, which I'm sure you know is often convenient in QFT. Since the time-domain is now finite, momentum-space integrals are replaced by Fourier series for the time-component, discretising the zero-component of the momenta: For bosons with a four-momentum $P = (p_0, \boldsymbol{p})$,

$$\int_{\mathbb{R}^d} \frac{\mathrm{d}^D P}{(2\pi)^D} f(P) \longrightarrow T \sum_{n=-\infty}^{\infty} \int_{\mathbb{R}^d} \frac{\mathrm{d}^d \boldsymbol{p}}{(2\pi)^d} f(2n\pi T, \boldsymbol{p}) \equiv \sum\!\!\!\!\!\!\int_P f(P).$$

There is nothing magical about this: It's simply a matter of Fourier analysis, and the fact that on finite intervals Fourier transforms become Fourier series (and usually the latter are even the starting point!). The discretised zero momenta are called *Matsubara modes*, the sums over them Matsubara sums, and even the whole imaginary-time formalism is sometimes known as the Matsubara formalism. The modes $p_0 = 2\pi n T$, $n \in \mathbb{Z}$ arise for bosons. For fermions, imposing antiperiodic boundary conditions for the Matsubara sum sets the condition $p_0 = (2n+1)\pi T$, $n \in \mathbb{Z}$. Fermionic Matsubara sums are oftne notationally associated with brackets: $\sum_{\{p_0\}}$ and $\sum\!\!\!\!\!\!\int_{\{P\}}$.

As a side note, Aleksi will talk about imaginary-time formalism at finite *density* later today, there the only change is that the zero-component of fermions (or charged scalars) is shifted by a constant factor $i\mu$, where μ is the chemical potential of that particle species. Often, this also goes under the thermal field theory umbrella.

Standard methods are valid for the spatial integrals, but computing Matsubara sums becomes a major complication for performing finite-temperature calculations. There are a variety of methods for dealing with such sums. One of them shows a close connection to statistical physics: By considering a function $\nu_B : x \mapsto i n_B(ix)$, where n_B is the Bose distribution $n_B(x) = \left(e^{\beta x} - 1\right)^{-1}$, and assuming that f is sufficiently regular we can write the Matsubara sum as a contour integral with

$$T \sum_{p_0} f(p_0, \boldsymbol{p}) = \sum_{p_0} \operatorname{Res}\left[f(x, \boldsymbol{p})\, \nu_B(x), x = p_0\right] = \int_{\Gamma_\delta} f(x, \boldsymbol{p})\, \nu_B(x) \frac{\mathrm{d}x}{2\pi i},$$

where $\Gamma_\delta = \cup_{p_0} C_\delta(p_0)$ and $\delta \in \mathbb{R}_+$ is chosen to be small enough that each circle $C_\delta(p_0)$ only contains a single Matsubara mode. Further deforming the contour leads to the representation

$$= \int_{\mathbb{R}} f(x, \boldsymbol{p}) \frac{\mathrm{d}x}{2\pi} + \int_{-\infty - i\varepsilon}^{\infty + i\varepsilon} \left[f(x, \boldsymbol{p}) + f(-x, \boldsymbol{p})\right] n_B(ix) \frac{\mathrm{d}x}{2\pi i},$$

where $\varepsilon \in \mathbb{R}_+$ and the first term represents zero-temperature contributions.

Another typical feature of thermal Matsubara sums is that in dimensional regularisation, they often lead to (Riemann or Hurwitz) ζ functions in the same way that zero-temperature calculations lead to Euler's Γ and B functions. We can see this with an explicit example: The sum-integral of the nth power of a scalar propagator leads to

$$\sum\!\!\!\!\!\!\int_P \frac{1}{P^{2n}} = \frac{\lambda\left(S^{d-1}\right)}{(2\pi)^d} \left(\frac{e^{\gamma_E} \bar{\Lambda}^2}{4\pi}\right)^{\frac{d-3}{2}} \times T \sum_{p_0} \overbrace{\int_0^\infty \mathrm{d}p \frac{p^{d-1}}{\left(p_0^2 + p^2\right)^n}}^{\text{cf. } \int_P \left(P^2 + m^2\right)^{-n} =} = \frac{\lambda\left(S^{d-1}\right)}{(2\pi)^d (4\pi)^{\frac{d}{2}}} \left(\frac{e^{\gamma_E} \bar{\Lambda}^2}{4\pi}\right)^{\frac{d-3}{2}} T \sum_{p_0} \frac{\Gamma\left(n - \frac{d}{2}\right)}{\Gamma(n)} \frac{1}{|p_0|^{2n-d}}$$

$$= \left(\frac{e^{\gamma_E} \bar{\Lambda}^2}{4\pi T^2}\right)^{\frac{d-3}{2}} \frac{\pi^{3/2}}{2 (2\pi)^{2n}} \frac{\Gamma\left(n - \frac{d-3}{2}\right)}{\Gamma(n)} T^{4-2n} \zeta(2n - d).$$

Here, we not only immediately see the ζ function appear, but also see the ramifications of introducing an intrinsic scale T: The integral would simply vanish in dimensional regularisation at $T = 0$ as a scalefree integral.

2.3 Leading-order results

The previous section is really all we need to start computing things. As usual, to get some familiarity with the new tools, we start with a toy model: The oh-so-familiar bosonic $\lambda\phi^4$ scalar field theory, which we take to be massless to get closed-form results—and actually, start by considering the free pressure. It is simply

$$\Omega_b = -\ln Z = -\ln \int \mathscr{D}\phi \exp\left(-\frac{1}{2}\int_X \partial_\mu\phi\partial^\mu\phi\right) = -\ln \int \mathscr{D}\phi \exp\left(-\frac{1}{2}\langle\phi, -\Delta\phi\rangle\right)$$

where Δ is the Laplacian on the newly-compactified space. By performing the Gaussian integral, making use of $\ln\det A = \operatorname{Tr}\ln A$ and moving to Fourier-space (dropping an irrelevant overall volume constant—if you wanted to be careful, you could look at $\ln Z/(\beta V)$ in a finite volume, and then take the infinite-volume limit, but I assume you know how to do this from vacuum QFT) we get

$$= \frac{1}{2}\ln \operatorname{Det}(-\Delta) = \frac{1}{2}\operatorname{Tr}\ln(-\Delta) = \frac{1}{2}\sum\!\!\!\!\!\!\int_P \ln P^2.$$

This integral is often represented by a single loop, but the representation is misleading at this order. The sum-integral can be computed by writing, up to an irrelevant integration constnat,

$$\ln\left(p_0^2 + x^2\right) = \int \mathrm{d}x \frac{2x}{p_0^2 + x^2} \implies T\sum_{p_0} \ln P^2 = \int \mathrm{d}x \sum_{p_0} \frac{2xT}{p_0^2 + x^2}.$$

This is the most fundamental Matsubara sum. There are many ways to do this, and Mathematica will simply give you the result. We can for example use the n_B-integral representation from above, and do a contour deformation to pick up the correct residue, to obtain

$$= \int \mathrm{d}x \frac{2xT}{2xT}\coth\frac{x}{2T} = x + 4T\operatorname{arctanh}\left(1 - 2e^{x\beta}\right) \sim 2T\ln\left(1 - e^{x\beta}\right),$$

where $\sim$ indicates that all scalefree contributions (polynomial in T) were dropped. Computing the integral

$$\Omega_b = \frac{1}{2}\sum\!\!\!\!\!\!\int_P \ln P^2 = \frac{1}{4\pi^2}\int_0^\infty \mathrm{d}p\, p^2 2T\ln\left(1 - e^{x\beta}\right) = -\frac{\pi^2}{90}T^4.$$

This is, of course, the (minus) free pressure of a gas of bosons, as it should be — and we could've taken the result directly from statmech, but I think it's neat

It is not too difficult to compute the (Dirac) fermionic equivalent in free theory: The path-integral gives $\Omega_f = -\ln\operatorname{Det}(\not{\partial}) = -2\sum\!\!\!\!\!\!\int_{\{P\}} \ln P^2$ and in $d = 3$ the sum-integral can be done either directly or by noting that $\sum_{\{p_0\}} f[p_0(T)] = 2\sum_{p_0} f\left[p_0\left(\frac{T}{2}\right)\right] - \sum_{p_0} f[p_0(T)]$. In the end, we get $\sum\!\!\!\!\!\!\int_{\{P\}} \ln P^2 = \frac{7\pi^2}{360}T^4 = \frac{7}{8}\frac{\pi^2}{45}T^4$.

With these, we can obtain the free pressure of, say, $SU(N_c)$ Yang–Mills theory with N_f massless Dirac fermions [remember–the ghosts have a fermionic path integral, but a standard bosonic operator] in an R_ξ-gauge:

$$\Omega_{\mathrm{YM}}^{\mathrm{LO}} = -\ln \frac{\operatorname{Det}\not{\partial}^{N_c N_f}\operatorname{Det}(-\Delta)^{d_A}}{\operatorname{Det}\left(-\delta_\alpha^\beta\Delta + (1 - \xi^{-1})\partial^\beta\partial_\alpha\right)^{d_A}} = \frac{d_A}{2}\left(D\sum\!\!\!\!\!\!\int_P \ln P^2 - \sum\!\!\!\!\!\!\int_P \ln\xi\right) - 2N_c\sum\!\!\!\!\!\!\int_{\{P\}} \ln P^2 - d_A\sum\!\!\!\!\!\!\int \ln P^2$$

$$= -\frac{\pi^2}{45}\left(d_A + \frac{7}{4}N_c N_f\right)T^4!$$

The gauge-dependence cancels, but the ghosts were essential.

2.4 Loop corrections

Perturbative corrections work just like they do in vacuum: For the $\lambda\phi^4$ theory, we can write (including only connected graphs, which comes about by taking the log)

$$\Omega_{\lambda\phi^4} = \Omega_b + \langle S_I \rangle - \frac{1}{2}\left\langle S_I^2 \right\rangle + \ldots,$$

and perform all the Wick contractions, or we can simply draw the diagrams and compute with the usual Feynman rules, just with sum-integrals instead of integrals (and Kronecker deltas instead of Dirac ones):

$$\langle S_I \rangle = \frac{3}{4!}\lambda \left(\mathop{\sum\!\!\!\!\!\!\int}_P \frac{1}{P^2} \right)^2 = 3 \times$$ 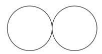 $$= \frac{\lambda T^4}{4^3 \times 3}$$

In the massless theory, this is finite. In a massive theory, there will be a mass-dependent divergence from a "vacuum×matter" term (for the LO pressure, we vacuum-subtract the divergence away), but it can be renormalised away in the usual way (note that the leading-order pressure must be similarly amended to be massive). Proceeding to next order,

$$-\frac{1}{2}\left\langle S_I^2 \right\rangle = -36 \times$$ 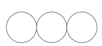 $$- 12 \times$$ 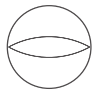 $$= -\frac{9}{144}\lambda^2 \left(\mathop{\sum\!\!\!\!\!\!\int}_P \frac{1}{P^2} \right)^2 \mathop{\sum\!\!\!\!\!\!\int}_Q \frac{1}{Q^4} - \frac{3}{144}\lambda^2 \mathop{\sum\!\!\!\!\!\!\int}_{PQR} \frac{1}{P^2 Q^2 (Q-R)^2 (R-P)^2}$$

The second integral is more tricky, but leads to a result $T^4 \left[e^{\gamma_E} \bar{\Lambda}^2 / (4\pi T)^2 \right]^{3\varepsilon} \left[\varepsilon^{-1} + \frac{91}{15} + 8\frac{\zeta'(-1)}{\zeta(-1)} - 2\frac{\zeta'(-3)}{\zeta(-3)} + O(\varepsilon) \right] / \left[24 (4\pi)^2 \right]$—again with an ultraviolet divergence that can be renormalised away. A very far-sighted reader might suspect that something is off in the first term because of the *infrared* divergent integral $\mathop{\sum\!\!\!\!\!\!\int}_Q \frac{1}{Q^4}$, but seemingly everything goes fine, and after renormalising we get

$$\frac{\Omega_{\lambda\phi^4}}{-T^4\pi^2/90} = 1 - \frac{5}{64\pi^2}\lambda + \frac{5}{4\left(64\pi^2\right)^2} \left[3\ln\frac{\bar{\Lambda}}{4\pi T} + \gamma_E + \frac{31}{15} + 4\frac{\zeta'(-1)}{\zeta(-1)} - 2\frac{\zeta'(-3)}{\zeta(-3)} \right] \lambda^2 + O\left(\lambda^4\right).$$

However, it turns out that something is missing: There is a $O\left(\lambda^{3/2}\right)$-term that gets added between the $O\left(\lambda^2\right)$ and $O(\lambda)$-terms, and we will return to that near the end. This is more obvious if one considers a massless theory, as there this missing feature will show up as a term $\propto m^{-1}$ in the small-mass limit already at $O\left(\lambda^2\right)$.

The power corrections can similarly be computed for the Yang–Mills theory: In Feynman gauge, everything factorises, and at $O\left(g_s^2\right)$ there are no problems: One gets

$$\Omega_{\text{QCD}}^{\text{NLO}} = \ldots = \frac{g_s^2}{4} d_A (d-1) \left\{ N_c (d-1) \left(\mathop{\sum\!\!\!\!\!\!\int}_P \frac{1}{P^2} \right)^2 - N_f \frac{\dim \mathcal{C}\ell(1,d)}{2} \left[2 \left(\mathop{\sum\!\!\!\!\!\!\int}_P \frac{1}{P^2} \right) \left(\mathop{\sum\!\!\!\!\!\!\int}_{\{Q\}} \frac{1}{Q^2} \right) - \left(\mathop{\sum\!\!\!\!\!\!\int}_{\{Q\}} \frac{1}{Q^2} \right)^2 \right] \right\} = d_A \left(N_c + \frac{5}{4} N_f \right) \frac{g_s^2 T^4}{144}!$$

The three-loop diagrams are a considerably more difficult undertaking, but in the end lead to a much more apparent problem than in $\lambda\phi^4$: There is an uncancelled IR-divergence. This is very characteristic of thermal field theory, causes a vast amount of problems, and will be the topic of the last part of the lectures, where introducing effective field theories lets one address it, as well as to add the missing part to the $\lambda\phi^4$.

Exercises

- (EXTRA) The derivation of the imaginary-time formalism was quite heuristic, and required the reader to believe certain definitions and associations. If you want to further convince yourself of its worth, directly obtain a path-integral of a quantum mechanical particle with a Hamiltonian $H = \frac{p^2}{2m} + V$ from $Z = \mathrm{Tr}e^{-\beta H}$. You should obtain $Z = \int_{x(\beta)=x(0)} \mathscr{D}x \exp\left\{-\int_0^\beta \mathrm{d}t \left[\frac{m}{2}x_t'' + V(x_t)\right]\right\}$, and see eg. the periodic boundaries arise in a natural way. The QFT result then follows by the same arguments as when constructing the path-integral formalism for QFT from QM. *Hint:* This is discussed extensively in practically all textbooks, eg. Laine & Vuorinen.
- (EXTRA) Confirm the contour representation of the Bosonic Matsubara sum by evaluating the residues, and derive an equivalent fermionic expression. If you are mathematically inclined, make sure you understand the conditions f (and do a "there exists $\delta > 0$" proof) must fulfill in the construction. Compute the standard bosonic Matsubara sum by hand, eg with a contour integral or a Fourier transform.
- By using standard rules of Grassmann integration and performing the logarithmic fermionic sum-integral, confirm that the formula stated in the main text for the pressure of a free fermionic field is correct.
- Using the QCD Feynman rules, check the contractions in the evaluation of the NLO pressure

3 Real-time formalism

3.1 Deriving the field-doubling

Before trying to tackle the infrared divergences, let us remind ourselves of the KMS condition from the first part:

$$\left\langle \mathcal{O}_1^M(t_1,\boldsymbol{x}_1)\,\mathcal{O}_2^M(t_2,\boldsymbol{x}_2)\right\rangle = \left\langle \mathcal{O}_2^M(t_2-i\beta,\boldsymbol{x}_1)\,\mathcal{O}_1^M(t_1,\boldsymbol{x}_2)\right\rangle .$$

This implies some form of periodicity in the coordinate must remain, but with real times t_1, t_2 this may not be complete. Considering simply a composite operator $\mathcal{O}(t)$, eg. here noting that the correlation function can only depend on $t_1 - t_2$ and writing as such $t = t_1 - t_2$ and writing $\mathcal{O}_1(t_1,\boldsymbol{x}_1)\,\mathcal{O}_2(t_2,\boldsymbol{x}_2) \equiv \mathcal{O}(t)$, and noting also that now the density operator is not necessarily $e^{-\beta H}$ as we allow small deviations from equilibrium, and assume that it can be written as $\rho(t) = \frac{1}{Z}e^{-iHt}\rho(t')\,e^{iHt}$. Now

$$\langle \mathcal{O}(t)\rangle = \int_{\Phi_i} \mathrm{d}\Phi_i \left\langle \Phi_i \left| \mathcal{O}(t)\, e^{-iHt}\rho(t')\, e^{iHt}\right| \Phi_i\right\rangle$$

$$= \int \mathrm{d}\Phi \mathrm{d}\Phi' \mathrm{d}\Phi'' \mathrm{d}\Phi''' \left\langle \Phi \left| \mathcal{O}(t) \right| \Phi' \right\rangle \left\langle \Phi' \right| e^{-iHt} \left| \Phi'' \right\rangle \left\langle \Phi'' \right| \rho(t') \left| \Phi''' \right\rangle \left\langle \Phi''' \left| e^{iHt} \right| \Phi \right\rangle$$

If we now choose t s.t. $\rho(t')$ is the equilibrium density matrix, and fix $t' = 0$ wlog,

$$= \frac{1}{Z}\int \mathrm{d}\Phi \mathrm{d}\Phi' \mathrm{d}\Phi'' \mathrm{d}\Phi''' \int_{\Phi_E(0)=\Phi'''}^{\Phi_E(-i\beta)=\pm\Phi''} \mathscr{D}\Phi_E \int_{\Phi_1(0)=\Phi''}^{\Phi_1(t)=\Phi'} \mathscr{D}\Phi_1 \int_{\Phi_2(0)=\Phi'''}^{\Phi_2(t)=\Phi} \mathscr{D}\Phi_2 \left\langle \Phi \left| \mathcal{O}(t) \right| \Phi' \right\rangle \exp\left(iS[\Phi_1] - iS[\Phi_2] - S_E[\Phi_E]\right).$$

We see that the field has *doubled* (tripled, really, but we may identify Φ_2 and Φ_E via an analytic continuation due to the matching boundaries in origin). The usual interpretation of this is that the fields Φ_1 travel between $(0,t)$ and then fields Φ_2 move first from $(0,t)$ and then over $(0,-i\beta)$. One could then write eg.

$$\langle \mathcal{O}(t)\rangle = \frac{1}{Z}\int \mathscr{D}\Phi_1\Phi_2\Phi_E \mathcal{O} \exp\left(iS[\Phi_1] - iS[\Phi_2] - S_E[\Phi_E]\right) = \frac{1}{Z}\int \mathscr{D}_{\mathcal{C}}\Phi_{\mathcal{C}} \mathcal{O} \exp\left(iS[\Phi_{\mathcal{C}}]\right),$$

where the implication is that the fields move along this Schwinger–Keldysh contour with the above boundary conditions (Which simply boil down to (anti)periodicity on the thermal circle, agreement at the origin and t). The geometrical identification of this is that the time factor of the base space becomes a cylinder or a Möbius strip realised as a submanifold of $\mathbb{C}$ with dimension two.

Just like the imaginary-time formulation, this has immediate consequences: For example, all propagators $\langle\phi(X)\,\phi^*(Y)\rangle_{\phi\in\Phi}$ are now 2×2-matrices in "Schwinger–Keldysh indices", and more generally n-point functions are n-dimensional arrays. Let us elucidate this with an example: The free gluon propagator

$$G^{ab}_{\mu\nu}(X,Y) = \left\langle \mathrm{T} A^a_\mu(X)\, A^b_\nu(Y)\right\rangle \mapsto_{ij} \mathbb{G}^{ab}_{\mu\nu}(X,Y) = \left\langle \mathrm{T}_i A^a_\mu(X)_j\, A^b_\nu(Y)\right\rangle = \begin{pmatrix} \left\langle {}_1A^a_\mu(K)_1\, A^b_\nu(P)\right\rangle & \left\langle {}_1A^a_\mu(K)_2\, A^b_\nu(P)\right\rangle \\ \left\langle {}_2A^a_\mu(P)_1\, A^b_\nu(K)\right\rangle & \left\langle {}_2A^a_\mu(K)_2\, A^b_\nu(P)\right\rangle \end{pmatrix}.$$

where the indices i,j indicate the position of the gluon fields on the Schwinger–Keldysh contour[4] and T denotes time ordering, implicitly along the Schwinger–Keldysh contour when applicable. The propagator can be derived directly from the generating functional as usual, but this is quite a tedious process. Alternatively, it can be analytically continued in an appropriate domain, but this requires defining about a dozen distinct Green's functions with differing causal properties, which the KMS relation will connect to each other. The result can be written in a form that turns out to contain the distribution function n_B , showing that it now appears explicitly in the propagator. Instead of doing that, we use a slightly more convenient basis.

3.2 The r/a-basis

The 1/2-fields can of course be transformed in various ways. The so-called r/a-basis (popularised by Caron-Huot) is something that often comes in handy especially with loop computations, as it allows for a convenient diagrammatic expansion. It is defined by the change of basis associated with the following linear transformation :

$$(\phi_1,\phi_2) \mapsto \mathcal{C}(\phi_1,\phi_2) \equiv (\phi_r,\phi_a)\,,\ \mathcal{C} = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ 1 & -1 \end{pmatrix};\ \mathcal{C}^{-1} = \begin{pmatrix} 1 & \frac{1}{2} \\ -\frac{1}{2} & 1 \end{pmatrix}.$$

We can finally write down the free gluon propagator

$$\begin{aligned} \tilde{\mathbb{G}}^{ab}_{\mu\nu}(K,P) &= \delta^{ab}\delta(K-P)\left\{ g_{\mu\nu}\begin{pmatrix} \left[\frac{1}{2}+n_B(p^0)\right]2\pi\mathrm{sgn}(p_0)\,\delta(P^2) & -i\left(P^2-i\eta p^0\right)^{-1} \\ -i\left(P^2+i\eta p^0\right)^{-1} & 0 \end{pmatrix} - (1-\xi)\,P_\mu P_\nu \begin{pmatrix} \left[\frac{1}{2}+n_B(p^0)\right]\left[2\pi\mathrm{sgn}(p_0)\,\delta(P^2)\right]^2 & \left(P^2-i\eta p^0\right)^{-2} \\ \left(P^2+i\eta p^0\right)^{-2} & 0 \end{pmatrix}\right\} \\ &\equiv \begin{pmatrix} D^{rr} & D^R \\ D^A & 0 \end{pmatrix}^{ab}_{\mu\nu}, \end{aligned}$$

Here R,A refer to retarded and advanced causal propagators and D^{rr} is the so-called symmetric propagator, $\eta\in\mathbb{R}_+$, and we have identified

$$-i\lim_{\eta\to 0}\left[\left(P^2-i\eta p^0\right)^{-1} - \left(P^2+i\eta p^0\right)^{-1}\right] = 2\pi\mathrm{sgn}(p_0)\,\delta\left(P^2\right)$$

and similarly for the powers of δ, in a distributional sense.

To discuss the diagrammatic consequences, let's consider $\lambda\phi^4$ instead for simplicity. The propagator transforms in an analogous way, with three nonzero components. The *interaction* transforms from $\frac{1}{4!}\lambda\phi_1^4 - \frac{1}{4!}\lambda\phi_2^4$ to $\frac{1}{4!}\left(\lambda\phi_r^3\phi_a + \frac{1}{4}\lambda\phi_a^3\phi_r\right)$: The number of vertices remains the same, but in the r/a-basis there is a factor of $\frac{1}{4}$ in one of them. This turns out to happen for QCD as well. Diagrammatically, we can associate a line from an a-field to an r-field with the R-propagator, from an r-field to an a-field with a A-propagator, and from an r-field to an r-field with an rr-propagator, with the aa-propagator always vanishing, *and draw causal flow arrows which always*

[4] This is not standard notation, but I hate having adjacent indices that belong to different spaces.

move towards the r-fields. Likewise, $\phi_r^3\phi_a$ would correspond to an interaction with three incoming and one outgoing causal arrow, and vice versa for $\phi_a^3\phi_r$. Then, all the diagrams generated for a given N-point function will be such that at L loops they will contain exactly L rr-propagators, and must have continuous causal flow from the beginning to the end. For example, the one-loop retarded self-energy term in $\lambda\phi^4$ arises from a single r/a-basis diagram:

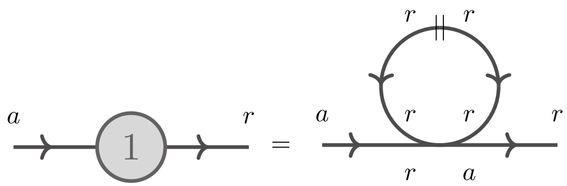

Since none of the $1/2$-basis propagators vanish, both the ϕ_1^4-vertex and the ϕ_2^4-vertex diagrams must be considered in that basis. This becomes increasingly useful at higher orders.

Exercises

- Check the r/a change-of-basis-rules explicitly.
- In the $\lambda\phi^4$ and the real-time formalism, write down the r/a-basis diagrams for the two-loop corrections to the retarted self-energy $\Pi^R(P)$.

4 EXTRA Thermal EFTs

At this point, we can go back to the leftover infrared divergence that left us puzzled with the perturbative calculation, and the missing piece I claim was absent from $\lambda\phi^4$. In both bosonic $\lambda\phi^4$ and QCD, they turn out to be associated with the integral $\oint_Q \frac{1}{Q^4}$[5]. Clearly, something goes wrong when $Q \to 0$. Looking at the integral more carefully, it turns out that all contributions for which $q_0 > 0$ are finite, and the potential problem is the *zero mode* $n = 0$.

4.1 Dimensional reduction in $\lambda\phi^4$

In EFT terms, this is the *soft* mode, and we can build a Wilsonian EFT by integrating out the *hard* modes for which $n \in \mathbb{Z}\setminus\{0\}$—these correspond to different physical scales: Starting with $\lambda\phi^4$, these are the soft scale $\lambda^{1/2}T$ and the hard scale T. Before constructing the EFT, let us consider how these scales contribute to the full pressure: The hard scale should contribute as $\Omega_b \sim \oint_P \ln P \sim T^4 n_B(T) \sim O(T^4)$, and the soft scale would contribute as $\Omega_b^{\text{soft}} \sim \oint_P \ln P \sim (gT)^4\, n_B(gT) \sim O(g^3T^4) > g^4T^4$. This is called a *Bose enhancement:* Due to the behaviour of n_B for small arguments $[n_B(\Lambda) \sim T/\Lambda]$, the soft thermal scale contributes at higher orders than one would naïvely expect.

To actually construct the EFT, we can Fourier-transform the fields $\phi = \sum_{n\in\mathbb{Z}\setminus\{0\}} e^{in\tau}\phi_n + \phi_0$, and integrate out the first sum. This takes care of all the time-dependence, so the EFT is a *three-dimensional*, *zero-temperature* theory with an action

$$S_{\text{Eff}} = \frac{1}{2}\partial_i\phi\partial^i\phi + \frac{1}{2}m_E^2\phi^2 + \frac{1}{4!}\lambda_E\phi^4,$$

which we have truncated to include only four-dimensional operators. Importantly, the theory has a mass even if the original theory is massless. The pressure is trivial to compute, as long as we remember that the new theory lives in $d = 3$, and yields (dropping all unimportant terms)

$$\Omega_{\text{Eff}} = -\frac{Tm_{E,R}^3}{12\pi} + \frac{T\lambda_{E,R}m_{E,R}^2}{8(4\pi)^2} + O\left(\lambda^5\right),$$

[5]More explicitly, the specific QCD diagram can be written as $\int \Pi^2(Q)\,Q^{-4}$ where the one-loop self-energy Π has a $\left[Q^2\right]^0$-term in the expansion near $Q^2 = 0$.

where we see the missing piece from the Ω_b is the pressure of a *three-dimensional massive theory at* $T = 0$. The matching coefficients can be computed in the usual way— we simply need $\lambda_{E,R} = \lambda_R T + O\left(\lambda^2\right)$, and $m^2_{E,R} = \frac{\lambda_R}{4!}T^2$, and we can obtain the proper pressure of the $\lambda\phi^4$ -theory.

$$\frac{\Omega_{\lambda\phi^4}}{-T^4\pi^2/90} = 1 - \frac{5}{64\pi^2}\lambda_R + \frac{5}{32\pi^3\sqrt{6}}\lambda_R^{3/2} + \frac{5}{4\left(64\pi^2\right)^2}\left[3\ln\frac{\bar{\Lambda}}{4\pi T} + \gamma_E + \frac{1471}{15} + 4\frac{\zeta'\left(-1\right)}{\zeta\left(-1\right)} - 2\frac{\zeta'\left(-3\right)}{\zeta\left(-3\right)}\right]\lambda^2 + O\left(\lambda^4\right).$$

It is interesting to note that in this case, it was not at all obvious that the EFT was necessary, since the naïve massless result *seemed* perfectly well-behaved at this order (note that at $O\left(\lambda^3\right)$, it would also start showing problems). The key point was identifying that there is a soft degree of freedom. It should also be noted that the soft scales attain a *thermal mass*: Physically, this corresponds to the well-known effect of Debye screening, where particles in a medium get slowed down by in-medium interactions—this is something that occurs even in nonrelativistic space plasmas!

4.2 Dimensional reduction in QCD

In QCD, it is obvious that something is wrong, but the structure of EFTs turns out to be much more complicated. The principle is the same: One identifies the hard degrees of fredom, and integrates them out. In this case, they are the quarks as a whole as well as the non-zero modes of the gauge fields. I am not going to go through the details, but the result is again a 3-dimensional theory, in this case a Yang–Mills theory with the A_0-field turning into a scalar in the adjoint representation. This theory is known as *electrostatic* QCD, or EQCD. It can be used to compute corrections to the theory without problems to $O\left(g^5\right)$. In Yang–Mills theory, this is not the end of the story, however: The scale $O\left(g^2T\right)$ is *also* distinct, and it starts to contribute at $O\left(g^6\right)$ in the small-coupling expansion. It can be inferred that now the softest scale, or *ultrasoft scale*, arises from the spatial gauge fields, and the temporal gauge field is now a harder (but still soft) component which can be integrated out. This results in a three-dimensional Yang–Mills theory called *magnetostatic* QCD (MQCD), but with a particularly annoying feature: It contributes nonperturbatively. The simplest argument is that the effective coupling at this scale, $g_s^2 n_B\left(g_s^2 T\right) = O\left(1\right)$, but this does not explain why this problem is not present in scalar theory. To do so, one can consider IR-contributions to the class of diagrams

It turns out that starting from four-loop order — at $O\left(g^6\right)$ — the structure of these diagrams is such that *every* diagram contributes to $O\left(g^6\right)$ from if the scales in the diagrams are $O\left(g^2T\right)$. This is a fundamental problem of finite-temperature perturbation theory known as the Linde problem, and it shows that one must obtain nonperturbative input (eg. from lattice) to proceed to high orders.

4.3 Hard thermal loops

The DR method is only really established in the imaginary-time formalism. However, curing the IR problems can be achieved also in a different way: By resumming. For example, resumming the problematic "daisy" graphs in $\lambda\phi^4$ in a geometric sum adds a mass-like term in the denominator of the integrand, which protects it in the infrared. The same effect happens in QCD, where this is easiest to see looking at the self-energy:

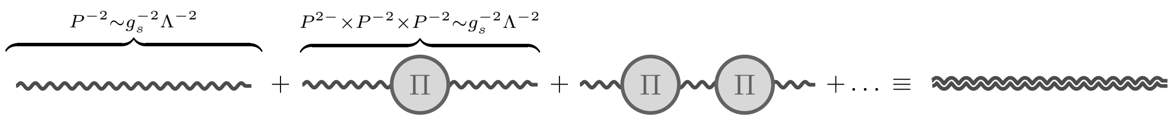

This is the standard Dyson resummation, and now the "problem" that $\Pi \to$ const at finite T becomes an advantage, since it protects the resummed propagator in the infrared. Such a resummation can also be performed in the real-time formalism without breaking the causal structure. Now, for example, the problematic QCD graph can be resummed to take care of the divergence:

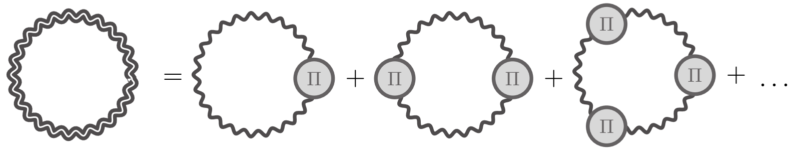

If one now considers the fact that the the theory only needs to be protected in the infrared, for soft momenta, it becomes obvious that in these resummed graphs the self-energy can be approximated by its soft limit, and the hard contributions can be computed naïvely (in practice, their IR divergences will cancel with the UV divergences of the resummed and soft-approximated graphs).

This idea leads to the EFT of hard thermal loops (HTL). It can also be motivated in a number of other ways — with a somewhat ambiguous Lagrangian, or by looking at Vlasov equations in non-Abelian plasma. However, I personally find this argument the simplest. It turns out that modifying the propagator is not enough, and the soft limits of loop corrections to other n-point functions must also be taken into account to guarantee the fulfillment of Ward identities, but this can be done, and leads to a consistent theory for treating soft gauge bosons and Dirac fermions. HTL computations are still under active consideration, but one of its significant downsides is that it is not a standard Wilsonian effective field theory. In its Lagrangian formulation, it is nonlocal, and only has a formal expansion parameter without a clear method for computing higher-order corrections to the Lagrangian.

Exercises

- Evaluate the pressure of the scalar DR theory to NLO.
- Evaluate the matching coefficient m_E^2 (and λ_E , if it is not obvious to you) in the scalar DR theory to LO.
- (A fair bit of work, but done in LV 8.4) Starting from the gluon self-energy

$$\Pi^{ab}_{\mu\nu}(K) = \frac{1}{2}\delta^{ab} g^2 N_c \sum\!\!\!\!\!\!\!\int_P \frac{1}{P^2 (K-P)^2}\left[\left(-4K^2 + 2(D-2)P^2\right)\delta_{\mu\nu} + (D+2)K_\mu K_\nu - 4(D-2)P_\mu P_\nu\right]$$

$$-\delta^{ab} g^2 N_f \sum\!\!\!\!\!\!\!\int_{\{P\}} \frac{1}{P^2 (K-P)^2}\left[\left(2P^2 - K^2\right)\delta_{\mu\nu} + 2K_\mu K_\nu - 4P_\mu P_\nu\right]$$

derive the HTL self-energy

$$H^{ab}_{\mu\nu}(K) = \frac{m_E^2}{2}\left[-\frac{k_0^2}{k^2} + \frac{ik_0}{2k}\left(1 + \frac{k_0^2}{k^2}\right)\ln\frac{ik_0 + k}{ik_0 - k}\right]\left(\delta_{ij} - \frac{k_i k_j}{k^2}\right)\delta_{\mu i}\delta_{\nu j}$$

$$+m_E^2\left[1 + \frac{k_0^2}{k^2}\right]\left[1 - \frac{ik_0}{2k}\ln\frac{ik_0 + k}{ik_0 - k}\right]\left[\delta_{\mu\nu} - \frac{K_\mu K_\nu}{K^2} - \left(\delta_{ij} - \frac{k_i k_j}{k^2}\right)\delta_{\mu i}\delta_{\nu j}\right]$$