INTRODUCTION TO QUANTUM FIELD THEORY

SHILIN LAI

This talk will be a miscellaneous selection of topics surrounding quantum field theory. We will first explain why it's needed over classical quantum mechanics. We then explain the operator/Heisenberg picture, doing some computations for the free field theory. To see relativistic invariance, we introduce the path integral/Lagrangian formulation. In $(1+1)$-dimensions, we will compute the theory for fields $S^1 \to S^1$ and recover T-duality. Finally, we reformulate everything in terms of functorial field theory, which clarifies the underlying algebraic structures.

The main references we used were Cambridge DAMTP lectures notes on QFT by David Tong and David Skinner, available on their websites. The canonical quantization part is an impressionistic sketch of the corresponding discussion in the Clay monograph *Mirror Symmetry*.

1. Why fields?

In classical Newtonian mechanics, a particle is described by a point in the phase space, and its motion is determined by a background potential. In the case of gravity, the motion of a particle leads to an instant change in this potential, which is problematic because we want *locality*, i.e. the behaviour of a system at a point should only depend on the information near a point. In the Lagrangian formalism, we want the action to be an integral of a Lagrangian which depends only on $x_i, \dot{x}_i, \ddot{x}_i, \ldots$, but not on something like $x_i - x_j$ since they take place at separate points in space. We resolve this non-locality in classical mechanics by postulating that long range forces are transferred at a finite speed via a *field*.

Example 1.1. We briefly recall classical electromagnetism. There are two fields: electric $\mathbf{E} : \mathbf{R}^3 \to \mathbf{R}$ and magnetic $\mathbf{B} : \mathbf{R}^3 \to \mathbf{R}^3$. In a vacuum, they satisfy the Maxwell equations

$$\nabla \cdot \mathbf{E} = 0 \qquad\qquad \nabla \cdot \mathbf{B} = 0$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \qquad\qquad \nabla \times \mathbf{B} = \frac{\partial \mathbf{E}}{\partial t}$$

In the presence of matter, we need to *couple* the equation by introducing terms on the right hand side. This system is clearly local, and the rearrange to give wave equations for $\mathbf{E}, \mathbf{B}$ with propagation speed of 1 (because we let $c = 1$).

Another reason to introduce fields in the quantum setting is to ensure relativistic invariance. The operator formulation of quantum mechanics fixes a specific time direction by singling out the Hamiltonian, so it is not relativistic. Instead, what we will do is to describe a state of the system as a field with both time and space parameters. An observable in this theory would be an operator valued field. To recover classical quantum mechanics, the expected value of the operator at $(\vec{x}, t)$ should be the probability density of finding the particle at time t at position $\vec{x}$.

We will first work out the commutation relation for the operators using a process called *canonical quantization* in the simple setting of *free fields*. It is not clear that the theory is relativistic, so we will later reformulate everything in terms of path integrals, which makes the propagator a central object instead of states and operators.

2. Canonical quantization

2.1. **General set-up.** We begin by setting up some physics terminologies. The background space-time for us is a $(1+d)$-dimensional manifold M. In physics, it has a Lorentzian metric with signature of $(-1, 1, \cdots, 1)$, and we write a vector as $(x^0, x^1, \cdots, x^d)$. Very soon, we will actually take M to be Riemannian by applying a *Wick rotation*. In physics, everything is written in coordinates, and being relativistic just means something is coordinate-independent.

A *field* will be a (sufficiently smooth) function $M \to$ (something), where the something can be $\mathbf{R}^n$ (bosonic/scalar fields), an exterior algebra (fermionic fields), a principal G-bundle (gauge fields), as well as other possibilities. We will focus on a single scalar field for now. Later, when we discuss T-duality, the target will be a Riemanniann circle.

The *Lagrangian* is a function $\mathcal{L}$ that takes a field and outputs a function on M. It should be *local* in the sense of the previous section, i.e. its value at a point only depends on the field and its derivatives at that point. The simplest example is the *free field Lagrangian*

$$\mathcal{L}_{\text{free}}(\phi) = -\frac{1}{2}\partial^\mu \phi \partial_\mu \phi$$

We can add a potential term $V(\phi)$. In a perturbative approach to calculations, a physicist would expand $V(\phi)$ into a Taylor series around 0. In this sense, the next simplest Lagrangian is

$$\mathcal{L}_{\text{KG}}(\phi) = -\frac{1}{2}\partial^\mu \phi \partial_\mu \phi - \frac{1}{2}m^2\phi^2$$

This is still a free field, which just means the equation of motion for ϕ is linear. The parameter m recovers the mass. Any higher powers of ϕ would lead to a non-linear equation of motion, and physicists call them interaction terms.

Given a Lagrangian, we can define the action to be the functional

$$\mathcal{S}[\phi] := \int_M \mathcal{L}(\phi) dx$$

The classical principle of least action is that the actual field is a critical point of the action. We can derive the Euler–Lagrange equation as usual. In particular, the Lagrangian $\mathcal{L}_{\text{KG}}$ leads to the equation

$$(\partial^\mu \partial_\mu + m^2)\phi = 0$$

This is the Klein–Gordon equation.

2.2. **Free field.** The picture so far is entirely classical. We now want to quantize a field described by the Klein–Gordon equation. For simplicity, we work in flat $(1+1)$-dimensions, so $M = \mathbf{R}_t \times \mathbf{R}_x$ with the Minkowski metric $\mathrm{diag}(-1, 1)$. The equation is explicitly written as

$$\left(-\partial_t^2 + \partial_x^2 + m^2\right)\phi = 0$$

This is a wave equation, which can be solved by several standard methods. Our approach will be to use the Fourier transform. Decompose ϕ in momentum modes

$$\phi = \frac{1}{2\pi}\int_{\mathbf{R}} e^{ipx}\phi(p,t)dp$$

then each $\phi(p,t)$ satisfies the equation $(-\partial_t^2 + (p^2+m^2))\phi(p,t) = 0$, which is a classical harmonic oscillator with $\omega_p = \sqrt{p^2+m^2}$. We recognize this as the relativistic energy of a particle of mass m and momentum p. Therefore, we decomposed he classical system as an infinite collection of harmonic oscillators with different frequencies. For quantization, we will use the quantum harmonic oscillators introduced in the previous talk.

For each $p \in \mathbf{R}$, we can define the Hilbert space $\mathcal{H}_p$ with creation and annihilation operators $a_p^\dagger, a_p$. They should satisfy $[a_p, a_p^\dagger] = 1$. The position, momentum, and Hamiltonian of the system are

$$\hat{x} = \frac{1}{\sqrt{2\omega_p}}(a_p + a_p^\dagger), \quad \hat{p} = -i\sqrt{\frac{\omega_p}{2}}(a_p - a_p^\dagger), \quad H_p = \hbar\omega_p\big(a_p^\dagger a_p + \frac{1}{2}\big)$$

so we can reconstruct the field of operators and its conjugate momentum by

$$\phi(x) = \frac{1}{2\pi}\int_{\mathbf{R}} \frac{1}{\sqrt{2\omega_p}}(a_p e^{ipx} + a_p^\dagger e^{-ipx})dp$$

$$\Pi(x) = \frac{1}{2\pi}\int_{\mathbf{R}} (-i)\sqrt{\frac{\omega_p}{2}}(a_p e^{ipx} - a_p^\dagger e^{-ipx})dp$$

They should satisfy the following *canonical commutation relations*

$$[\phi(x), \phi(y)] = [\Pi(x), \Pi(y)] = 0, \quad [\phi(x), \Pi(y)] = \delta(x-y)$$

After a computation, this reduces to the following

$$[a_p, a_q] = [a_p^\dagger, a_q^\dagger] = 0, \quad [a_p, a_q^\dagger] = 2\pi\delta(p-q)$$

The natural space of states should be the L^2-space over the moduli space of all fields. Having decomposed the field into Fourier modes, we can replace it by the *Fock space*

$$\mathcal{H}_{\text{Fock}} := \Big(\bigotimes_p{}' \mathcal{H}_p\Big)^\wedge$$

where the restricted tensor product means such that the right hand side is generated by pure tensors which is the vacuum state $|0\rangle$ at all but finitely many places, and the $(-)^\wedge$ means taking completion (more precisely, this is a direct integral of Hilbert spaces). Starting with the vacuum state $|0\rangle$, we can form states $|p_1 \cdots p_n\rangle := a_{p_1}^\dagger \cdots a_{p_n}^\dagger|0\rangle$. This state is physically n free particles travelling with momenta $p_1, \cdots, p_n$ respectively.

We can express the Hamiltonian in terms of creation and annihilation operators:

$$\begin{aligned} H &= \int_{\mathbf{R}} \frac{1}{2}\big(-\partial^\mu\phi\partial_\mu\phi + m^2\phi^2\big)dx \\ &= \frac{1}{4\pi}\int_{\mathbf{R}} \hbar\omega_p(a_p a_p^\dagger + a_p^\dagger a_p)dp \end{aligned}$$

There are two ways of seeing this. We can substitute in the expressions for ϕ and Π into the Hamiltonian, or we can view the Hamiltonian operator as the sum of the Hamiltonians over each individual harmonic oscillator. In any case, the vacuum energy is

$$H|0\rangle = \frac{1}{2}\int_{\mathbf{R}} \hbar\omega_p\delta(0)dp\,|0\rangle$$

This is infinity for two reasons. Firstly, $\delta(0) = \infty$. By Fourier inversion, $\frac{1}{2\pi}\int_{\mathbf{R}} e^{ipx}dp = \delta(x)$, so $\delta(0)$ is really the volume of space, which is infinite. This is known as *infrared divergence* because we are trying to sum up modes with low energy. This problem is usually dealt with by compatifying space. Upon replacing $\mathbf{R}_x$ by a flat circle of length L (i.e. imposing periodic boundary conditions), we resolve this infinity.

After compactifying space, the momenta are discretized to scales of L^{-1}, so we need to replace the integral with a sum. After rewriting m for mL, we get

$$H|0\rangle = \sum_{n=0}^{\infty}\sqrt{n^2+m^2}\,|0\rangle$$

This is still infinity, which is called *ultravioet divergence* because it happens at high energy modes. For some purposes, it is enough to say we only care about energy differences, which turns out to be equivalent to reordering the terms in the Hamiltonian into the *normal ordering* $a_p^\dagger a_p$.

We will not be satisfied with this because when we study T-duality, we need the whole partition function. One solution is to use zeta function regularization, which means we replace the sum with $\sum_n \sqrt{n^2+(m')^2}n^{-s}$ and evaluate at $s=0$ in the sense of analytic continuation. We will see a related computation later.

3. Periodic field and T-duality

Now we impose periodicity on the field values too, so a field is a smooth map $S^1 \times \mathbf{R}_t \to \mathbf{R}/2\pi R\mathbf{Z}$. For simplicity, we assume the source S^1 to have radius 1 and the field is massless. The simple example of T-duality we will study now is an equivalence of the system with parameter R and the system with parameter $\frac{1}{R}$. For now, this means we will identify some observables and the energy spectrum. Later, there will be a more robust reason why the two systems are physically equivalent.

Let x be the spatial coordinate, which is a variable in $\mathbf{R}/2\pi\mathbf{Z}$. The field variable ϕ takes values in $\mathbf{R}/2\pi R\mathbf{Z}$, and we will often view it as a function $\mathbf{R}_x \times \mathbf{R}_t \to \mathbf{R}$ such that $\phi(x+2\pi n, t) - \phi(x,t) \in 2\pi R\mathbf{Z}$. The action is derived from the usual free field Lagrangian,

$$\mathcal{S}[\phi] = \int_{\mathbf{R}/2\pi\mathbf{Z}} \frac{1}{2}\big(-(\partial_t\phi)^2 + (\partial_x\phi)^2\big)\frac{dx}{2\pi}$$

There are two conserved quantities for this system

$$p = \frac{1}{2\pi}\int_{\mathbf{R}/2\pi\mathbf{Z}} \partial_t\phi\,dx, \quad w = \frac{1}{2\pi}\int_{\mathbf{R}/2\pi\mathbf{Z}} \partial_x\phi\,dx \in R\mathbf{Z}$$

The winding number w is already "quantized", and the momentum will be quantized since we have a periodic boundary condition. The state space $\mathcal{H}$ should once again be L^2-functions on the space of all maps $S^1 \times \mathbf{R}_t \to \mathbf{R}/2\pi R\mathbf{Z}$. The analysis as in the previous section can be repeated to obtain a tensor product decomposition

$$\mathcal{H} = \mathcal{H}_{\text{top}} \otimes \mathcal{H}_0 \otimes \bigotimes_{n\neq 0}{}' \mathcal{H}_n$$

For $n > 0$, the space $\mathcal{H}_n$ is the state space for a harmonic oscillator with frequency $|n|$. For $n = 0$, the space $\mathcal{H}_0$ is not a harmonic oscillator but a free field valued in $\mathbf{R}/2\pi R\mathbf{Z}$. Mathematically, this theory is equivalent to a quantum particle in a periodic space, so it has a momentum operator $\hat{p}$ with eigenvalues in $\frac{1}{R}\mathbf{Z}$. Finally, we have a part $\mathcal{H}_{\text{top}}$ which describes the topology of ϕ. In our case it is decribed by the winding number. Therefore, we also have the decomposition

$$\mathcal{H} = \bigoplus_{k,l\in\mathbf{Z}} \mathcal{H}_{(k,l)}$$

where each $\mathcal{H}_{(k,l)}$ decomposes into a restricted tensor product of harmonic oscillators $\mathcal{H}_{(k,l),n}$, with operators $a_n, a_n^\dagger$ modifying particle numbers.

The spectrum of the system is usually packaged together into a function called the *partition function*

$$\mathcal{Z}(\beta) := \sum_E e^{-\beta E}$$

where the sum runs over all energy eigenvalues with multiplicity. When $\beta \in i\mathbf{R}_{>0}$, this can be viewed as the trace of time evolution, i.e.

$$\mathcal{Z}(it) = \operatorname{Tr}(e^{-itH})$$

The fact that we are using parameter β instead of t is another instance of Wick rotation. In functorial field theory, this will be the number attached to the torus with fundamental domain $[0, 2\pi] \times [0, t]$. Our goal is to compute $\mathcal{Z}(\beta)$ for our system. It is easy to check that the formation of $\mathcal{Z}(\beta)$ preserves algebraic operations, so we have

$$\mathcal{Z} = \sum_{l,m} \left(\mathcal{Z}[\mathcal{H}_{(l,m),0}] \prod_{n\neq 0} \mathcal{Z}[\mathcal{H}_{(l,m),n}]\right)$$

We can compute each term separately. All energies need to be multiplied by $\hbar$ which is absorbed into β

$$\mathcal{Z}[\mathcal{H}_{(l,m),n}] = \sum_{k=0}^{\infty} e^{-\beta|n|\left(k+\frac{1}{2}\right)} = e^{-\frac{|n|\beta}{2}} \frac{1}{1 - e^{-|n|\beta}}$$
$$\mathcal{Z}[\mathcal{H}_{(l,m),0}] = e^{-\frac{\beta}{2}((l/R)^2 + (mR)^2)}$$

When we take product over all $n \geq 1$, we get the divergent term $\prod_{n\geq 1} e^{-\frac{n\beta}{2}} = \exp(-\frac{\beta}{2}\sum_{n\geq 1} n)$, but famously, the sum is equal to $-\frac{1}{12}$, so

$$\mathcal{Z} = e^{-\frac{\beta}{12}} \prod_{n\neq 0} \frac{1}{1 - e^{-|n|\beta}} \sum_{(k,l)} e^{-\frac{\beta}{2}((l/R)^2 + (mR)^2)} = \frac{1}{|\eta(i\tau)|^2} \sum_{(k,l)} q^{\frac{1}{2}((l/R)^2 + (mR)^2)}$$

In the final expression, $\beta = 2\pi\tau$, $q = e^{-2\pi\tau}$, and η is the Dirichlet η-function. Recall that the partition function can be viewed as a number attached to a torus with fundamental domain $[0, 1] \times [0, \beta]$. We can change the structure of this torus by performing a space shift by τ' before gluing time slices $t = 0$ and $t = \beta$ together. The partition function for this torus is still given by the formula above, except now we are replacing $i\tau$ by $\tau' + i\tau$ everywhere. If we had not compactified the target, the partition function would just be the first part, which is a modular form in τ, i.e. it only depends on the conformal structure of the torus. Compactification introduces a length scale R, so it is natural that we see the size of the worldsheet too.

The key thing to note is that when $R \rightsquigarrow \frac{1}{R}$, the partition function remains the same, but the contribution from $\mathcal{H}_{(l,m)}$ is now coming from $\mathcal{H}_{(m,l)}$. Physically, we have the same theory, except we need to interchange certain sets of observables. .

4. Path integral formulation

The theory so far does not appear to be relativistic, since we fixed a decomposition of spacetime. Making it relativistic actually requires certain normalizing factors for a_p and $a_p^\dagger$. But it is relativistic since the observables of the theory are *propagators* $U(x,t;y,t')$ (or more generally *correlators*). This is a complex number whose norm is the probability density that a field concentrated at space-time coordinate (t,x) evolves to a field concentrated at (t',y). This can be computed and shown to satisfy the required properties.

Feynman's path integral approach makes relativity clear by making propagators the primary objects. The starting point is a reformulation of the action principle. Keeping the same general set-up, we can consider two arbitrary spacelike hypersurfaces M_0, M_1 and ask for the propagator between them. The path integral action principle states that the field takes all possibilities, weighed by the action

$$U(\phi_0,\phi_1) = \int_{\phi:\phi|_{M_i}=\phi_i} \exp\left(-\frac{i}{\hbar}\mathcal{S}[\phi]\right)\mathcal{D}\phi$$

We can view this as a linear operator from functions of fields on M_0 to functions on fields of M_1. As $\hbar \to 0$ in the above integral, only the critical points contribute by the method of stationary phase, so we recover the classical principle of least action in the classical limit. We can also write the partition function in a particularly simple form

$$\mathcal{Z} = \int \exp\left(-\frac{i}{\hbar}\mathcal{S}[\phi]\right)\mathcal{D}\phi$$

where the integral is over all fields on M. This is a number as opposed to a function between two function spaces of fields. If M is formed by taking its parts "between" M_0 and M_1 and gluing the end points together, the partition function is just the trace of the propagator. This is more clearly explained when we talk about functorial field theory. In particular, from this point of view, the Heisenberg picture only occurs on manifolds with boundaries.

The issue now is nothing we wrote down made sense. To solve this problem, we apply a *Wick rotation*: we change the time variable t to $i\tau$, making M into a Riemannian manifold. In other words, along the imaginary time axis, the transition amplitude is

$$\int_{\phi:\phi|_{M_i}=\phi_i} \exp\left(-\frac{\mathcal{S}[\phi]}{\hbar}\right)\mathcal{D}\phi$$

The factor of i disappears since we integrated over time for $\mathcal{S}$. The Lagrangian has a quadratic term which is now positive definite, so we get exponential decay away from the classical path. This is much better than cancellation coming from rapid oscillation. People who develop stochastic calculus can define the *Wiener measure*, which is essentially the non-existent path integral measure weighted by a Gaussian. Using analytic continuation, we can recover the supposed value of the oscillatory integral. From now on, we will do everything with a Riemannian metric.

One further subtlety is that the action involves derivatives of ϕ, but the Wiener measure is supported on nowhere differentiable paths. Let's not worry about it, but this is how path integral recovers non-commutativity in the operator picture.

To illustrate how to formally work with them, we re-derive T-duality where the spacetime is a general Riemann surface Σ with metric $g_{\mu\nu}$. Write $\phi = R\varphi$, then φ is periodic with period 2π, and the action is

$$\mathcal{S}_R[\varphi] = \frac{R^2}{4\pi}\int_\Sigma \partial^\mu\varphi\partial_\mu\varphi d^2\sigma$$

Here $d^2\sigma = \sqrt{g}dx^\mu dx^\nu$ is the volume form. We will define a new theory containing three fields which specialize to the above action for R and $\frac{1}{R}$ by taking the path integrals in different orders. To explain this point physically, suppose we are starting with a theory with fields φ, φ', but we are only interested in observations involving φ, we can perform an integration over φ' first and define an *effective action*

$$\mathcal{S}_{\text{eff}}[\varphi] = -\log\int \mathcal{S}[\varphi,\varphi']\mathcal{D}\varphi'$$

The two theories are identical from our point of view since we can't observe φ'. This is a fundamental process in physics to derive observable consequences from high energy theories.

When physicists derive T-duality, they consider the following theory with three fields: φ, θ are two periodic fields valued in $\mathbf{R}/2\pi\mathbf{Z}$, and $A = (A_1, A_2)$ is a *gauge field* valued in $\mathbf{R}^2$

$$\mathcal{S}[\varphi, \theta, A] = \frac{R^2}{4\pi}\int_\Sigma (\partial^\mu\varphi + A^\mu)(\partial_\mu\varphi + A_\mu)d^2\sigma - \frac{i}{2\pi}\int_\Sigma A_\mu\epsilon^{\mu\nu}\partial_\nu\theta d^2\sigma$$

where $\epsilon^{\mu\nu}$ is the tensor $\epsilon^{01} = 1, \epsilon^{10} = -1, \epsilon^{00} = \epsilon^{11} = 0$. In mathematical language, the field A introduces the 1-form $A = A_\mu dx^\mu$, giving rise to a connection on the tangent bundle. The first term is then the inner product of $\nabla_A\varphi$ with itself. The second term is the integral of $A \wedge d\theta$, which is a Lagrange multiplier to "fix the gauge". To see what this means, we perform the path integral $\int \mathcal{D}\theta$. This again requires keeping track of topological invariants. Write

$$d\theta = d\theta_0 + \sum_{i=1}^{2g} 2\pi n_i\omega^i$$

where $\{\omega^1, \cdots, \omega^{2g}\} \subseteq H^1_{\mathrm{dR}}(\Sigma, \mathbf{R})$ is the dual basis to a basis $\{\gamma_1, \cdots, \gamma_{2g}\}$ of $H_1(\Sigma, \mathbf{Z})$. The coefficients n_i describe the induced map $\theta_* : H_1(\Sigma, \mathbf{Z}) \to H_1(\mathbf{R}/2\pi\mathbf{Z}, \mathbf{Z})$. They are all integers.

Now the second term of the action is $\frac{i}{2\pi}$ times

$$\int_\Sigma A \wedge \left(d\theta_0 + \sum_{i=1}^{2g} 2\pi n_i\omega^i\right) = -\int_\Sigma \theta_0 dA + \sum_{i=1}^{2g} 2\pi n_i \int_\Sigma A \wedge \omega^i$$

First perform the path integral over all θ_0, which is now over all functions $\Sigma \to \mathbf{R}$

$$\int \exp\left(-\int_\Sigma \theta_0 dA\right)\mathcal{D}\theta_0$$

By Laplace's method, this is a measure concentrated at $dA = 0$, so what we have done is impose this condition in the original theory. We can write $A = df + \sum_{i=1}^{2g} a_i\omega_i$, where f is a function and $a_i \in \mathbf{R}$. The remaining action is now

$$\frac{i}{2\pi}\sum_{1\le i,j\le 2g} 2\pi n_j a_i \int_\Sigma \omega^i \wedge \omega^j$$

The "path integral" over the remaining modes is just a sum over all $n_i \in \mathbf{Z}$. Symbolically, one such sum looks like

$$\sum_{n\in\mathbf{Z}} e^{-ina} = 2\pi\sum_{m\in\mathbf{Z}} \delta(a - 2\pi m)$$

Therefore, we have the further constraint that $a_i \in 2\pi\mathbf{Z}$, or equivalently, $A = d\psi$ for a periodic field ψ. This can be absorbed into φ to recover the original action. Now, the integral over all fields A is just multiplication by the "volume" of the space of all fields. Whatever this means, it can be normalized in the end. Therefore, we have shown that $\mathcal{S}_{\mathrm{eff},\theta,A}[\varphi] = \mathcal{S}_R[\varphi]$.

On the other hand, we can integrate out φ first and calculate the effective action for θ and A. Note that a change of variable $\varphi \mapsto \varphi + \lambda$ and $A_\mu \mapsto A_\mu - \partial_\mu\lambda$ leaves the theory invariant, so the integral over φ is equal to setting $\varphi = 0$ and multiplying by a certain normalizing volume term. Now,

$$\begin{aligned}\mathcal{S}_{\mathrm{eff},\varphi}[\theta, A] &= \frac{R^2}{4\pi}\int_\Sigma A^\mu A_\mu d^2\sigma - \frac{i}{2\pi}\int_\Sigma A_\nu\epsilon^{\mu\nu}\partial_\mu\theta d^2\sigma \\ &= \frac{R^2}{4\pi}\int_\Sigma \left(A^\mu A_\mu - \frac{2i}{R^2}A_\nu\epsilon^{\mu\nu}\partial_\mu\theta\right)d^2\sigma\end{aligned}$$

The action has a critical point at $A^\mu = \frac{i}{R^2}\epsilon^{\mu\nu}\partial_\nu\theta$. The classical method of stationary phase is exact for the Gaussian integral, so we assume that it is also exact. Therefore, up to a normalizing constant, we just need to substitute the critical point into the expression to obtain the effective action

$$\mathcal{S}_{\mathrm{eff},\varphi,A}[\theta] = \frac{1}{4\pi R^2}\int_\Sigma \partial^\mu\theta\partial_\mu\theta d^2\sigma = \mathcal{S}_{\frac{1}{R}}[\theta]$$

which is the action for the σ-model with a target circle of radius $\frac{1}{R}$. This recovers a strong form of T-duality. Using the relations we have derived between A, θ, φ, one can also show that this duality interchanges momentum and winding numbers, as seen before.

5. Functorial field theory

We have seen that it is difficult to make the notion of path integrals rigorous. To take mathematical advantage of this notion, we will forget the analytic difficulties by axiomatizing the construction.

To start with, suppose we have a manifold M of dimension $d+1$ with boundary $\partial M = M_0 \sqcup M_1$. The propagator was something sending field configurations on M_0 to field configurations on M_1, so we have the following assignments

$$M_i^d \rightsquigarrow \mathcal{H}_i,\ i = 0, 1, \quad M^{d+1} \rightsquigarrow f \in \mathrm{Hom}(\mathcal{H}_0, \mathcal{H}_1)$$

If we have another $(d+1)$-dimensional manifold M' with boundary $M_1 \sqcup M_2$, then they can be glued to form $M \cup_{M_1} M'$ with boundary $M_0 \sqcup M_2$. The propagator should just be a composition of the two propagators, so we have a functor

$$\mathbf{Bord}^d \to \mathbf{Vect},\ M^d \rightsquigarrow \text{(states)},\ M^{d+1} \rightsquigarrow \text{(propagator)}$$

Definition 5.1. The d-borsim category $\mathbf{Bord}^d$ is the category with objects given by oriented d-manifolds (together with $\emptyset$) and morphisms $\mathrm{Mor}(M_0, M_1)$ given by $(d+1)$-manifolds with boundary $M_0 \sqcup (-M_1)$. It has a tensor product given by disjoint union making it a symmetric monoidal category.

We have also seen that physicists like to attach a number called the partition function to a closed $(d+1)$-dimensional manifold. To see this in our current point of view, the functor should attach the empty d-manifold to the 1-dimensional space $\mathbf{C}$, so a closed $(d+1)$-manifold is a function $\mathbf{C} \to \mathbf{C}$, which is a single number. This can be also seen from a cut-and-paste construction, as we have seen in the case of the torus in our discussion of T-duality. In the functorial formulation, this defines a function $\mathrm{Mor}(\emptyset, \emptyset) \to \mathbf{C}$, so what we are actually doing is exploiting the monoidal structure of $\mathbf{Bord}^d$.

Definition 5.2. A *quantum field theory* in dimension d is a symmetric monoidal functor from the d-bordism category (with extra structure) to the category of vector spaces.

This is enough for *topological quantum field theory*, where the action does not depend on the metric. Sometimes we want the theory to also depend on additional structures of M such as the metric structure, as we have seen in the path integral derivation of T-duality. This is reflected by putting the appropriate structures on objects in $\mathbf{Bord}^d$. For example, T-duality is the equivalence of two QFTs on the 1-bordism category of Riemannian manifolds.

The definition can be definitionally expanded into a list of axioms describing how the state space and partition functions behave with respect to disjoint union operations. They were formulated in explicit terms by Atiyah. The same paper also contains more discussions on the consequences of the axiom.