Workshop on Quantum Field Theory and Topological Phases via Homotopy Theory and Operator Algebras

Summer 2025

Program abstract: Quantum Field Theory (QFT) and Quantum Statistical Mechanics are central to high energy physics and condensed matter physics; they also raise deep questions in mathematics. The application of operator algebras to these areas of physics is well-known. Recent developments indicate that to understand some aspects QFT properly a further ingredient is needed: homotopy theory and infinity-categories. One such development is the recognition that symmetry in a QFT is better described by a homotopy type rather than a group (so-called generalized symmetries). Another one is the work of Lurie and others on extended Topological Field Theory (TFT) and the Baez-Dolan cobordism hypothesis. Finally, there is a conjecture of Kitaev that invertible phases of matter are classified by homotopy groups of an Omega-spectrum. This workshop will bring together researchers and students approaching this physics using different mathematical techniques: operator algebras, homotopy theory, higher category theory, etc. The goal is to catalyze new interactions between different communities. At the workshop recent developments will be reviewed and hopefully progress can be made on two outstanding problems: the Kitaev conjecture as well as the long-standing goal of finding a proper mathematical formulation for QFT.

This is an unofficial set of notes scribed by Gary Hu for the first week of lectures only. He is responsible for all mistakes. If you find any errors, please report them to: gh7@williams.edu

Contents

1 Bruno Nachtergaele: Ground States of Quantum Lattice Systems

Contents

1.1 Lecture 1

1.1.1 Introduction to Quantum Lattice States of Matter

These lectures will treat the subject of quantum lattice states of matter from a mathematical perspective. However, it's crucial to remember the physical motivation behind this study. We are driven by the existence of real-world materials and engineered systems that exhibit remarkable phenomena rooted in topology and quantum mechanics. At very low temperatures, matter can exist in highly entangled states, which dictate their response to external probes, their transport properties, and more. This field also forms the foundation of quantum information science, where the goal is to build devices that leverage these entangled states.

While we won't dive deeply into the physics, it's important to understand the origin of lattice systems. These are many-body systems with several key components:

1. **Atoms (Spins/Qudits):** These are the fundamental units, often possessing internal degrees of freedom called **spins** or, more generally, **qudits**. We typically model them as being located at regular positions in space, which we represent with a lattice or, more abstractly, a graph.
2. **Fermions (e.g., Electrons):** In addition to fixed atoms, there can be mobile particles, like electrons, that move between atomic sites.
3. **Phonons:** The locations of the atoms are not truly fixed. The vibrations of the crystal lattice are quantized into bosonic particles called **phonons**. They can play a role in mediating interactions between spins and electrons.

Our primary focus will be on **spin systems**, where the Hilbert space associated with each site is finite-dimensional. We will stick to the basics and the general framework for describing these important problems, assuming a foundational knowledge of quantum mechanics (Hilbert spaces, Banach spaces, etc.).

1.1.2 The General Setup: Lattices and Observables

The Lattice

The mathematical setting for our quantum spin systems begins with the "space" where our spins reside. We will take our lattice, Γ, to be a **discrete metric space** (Γ, d). This means that for any point $x \in \Gamma$ and any radius $r > 0$, the ball of radius r centered at x, defined as $B_x(r) := \{y \in \Gamma \mid d(x, y) \leq r\}$, contains a finite number of points.

We often assume additional properties for Γ.

Definition 1.1 (ν-regularity)**.** *A discrete metric space (Γ, d) is **ν-regular** if there exist constants $\kappa, \nu > 0$ such that the number of points in a ball of radius*

n grows at most polynomially:

$$|B_x(n)| \leq \kappa n^\nu.$$

for all $n \geq 1$.

Definition 1.2 (Delone Set)**.** *A set* $\Gamma \subset \mathbb{R}^\nu$ *is a* ***Delone set*** *if it is both uniformly discrete and relatively dense. This implies the existence of two constants:*

- *The* ***minimal distance****:* $\epsilon := \inf_{x \neq y \in \Gamma} d(x, y) > 0$.
- *The* ***covering radius****:* $\Delta := \sup_{z \in \mathbb{R}^\nu} \inf_{y \in \Gamma} d(z, y) < \infty$.

Often, for simplicity, results are formulated for the regular integer lattice $\Gamma = \mathbb{Z}^\nu$, but generalization to spaces with these properties is usually straightforward.

Observables

With the space defined, we can describe the quantum mechanical observables.

Definition 1.3 (Single-Site Algebra)**.** *For each site* $x \in \Gamma$*, the associated observables form the algebra* $\mathcal{A}_x$*, which is a copy of the* $d_x \times d_x$ *complex matrices,* $M_{d_x}(\mathbb{C})$*. The corresponding Hilbert space is* $\mathcal{H}_x \cong \mathbb{C}^{d_x}$.

Definition 1.4 (Local Algebra)**.** *For any* ***finite subset*** $\Lambda \subset \Gamma$*, the algebra of observables is the tensor product of the single-site algebras:*

$$\mathcal{A}_\Lambda = \bigotimes_{x \in \Lambda} \mathcal{A}_x$$

This algebra acts on the Hilbert space $\mathcal{H}_\Lambda = \bigotimes_{x \in \Lambda} \mathcal{H}_x$.

Definition 1.5 (Algebra of Local Observables)**.** *The algebra of* ***strictly local observables*** *is the union over all finite subsets of* Γ*:*

$$\mathcal{A}_{local} = \bigcup_{\Lambda \subset \Gamma, |\Lambda| < \infty} \mathcal{A}_\Lambda$$

*This is a norm-dense *-subalgebra, but it is not complete. For any two finite sets* $\Lambda_0 \subset \Lambda_1$*, the algebra* $\mathcal{A}_{\Lambda_0}$ *is embedded into* $\mathcal{A}_{\Lambda_1}$ *via the map* $A \mapsto A \otimes \mathbf{1}_{\Lambda_1 \setminus \Lambda_0}$.

Definition 1.6 (Quasi-Local Algebra)**.** *The* ***quasi-local algebra****,* $\mathcal{A}_\Gamma$*, is the C*-norm completion of the algebra of local observables:*

$$\mathcal{A}_\Gamma := \overline{\mathcal{A}_{local}}^{\|\cdot\|}$$

By definition, any observable $A \in \mathcal{A}_\Gamma$ *can be approximated by a sequence of local observables. That is, for any* $A \in \mathcal{A}_\Gamma$ *and any sequence of finite sets* Λ_n *that grows to encompass the whole lattice (*$\Lambda_n \nearrow \Gamma$*), there exists a sequence of local operators* $A_n \in \mathcal{A}_{\Lambda_n}$ *such that:*

$$\lim_{n \to \infty} \|A - A_n\| = 0$$

1.1.3 Dynamics of Infinite Systems

Interactions and Hamiltonians

The interesting physics arises from interactions between the spins.

Definition 1.7 (Interaction)**.** *An **interaction**, Φ, is a map that assigns a self-adjoint local observable to each finite subset of the lattice:*

$$\Phi : \{\textit{finite } X \subset \Gamma\} \to \mathcal{A}_{local}$$

such that for each finite set X, the interaction term $\Phi(X)$ is an element of $\mathcal{A}_X$ and is self-adjoint, $\Phi(X)^* = \Phi(X)$.

For any finite region $\Lambda \subset \Gamma$, we can now define the **local Hamiltonian** by summing all interaction terms supported within that region:

$$H_\Lambda = \sum_{X \subset \Lambda} \Phi(X)$$

This Hamiltonian generates the time evolution (Heisenberg dynamics) for observables within the finite system Λ:

$$\tau_t^\Lambda(A) = e^{itH_\Lambda} A e^{-itH_\Lambda}$$

for $A \in \mathcal{A}_\Lambda$.

The Thermodynamic Limit

To define the dynamics on the infinite lattice, we must take the **thermodynamic limit** ($\Lambda \to \Gamma$). This limit doesn't exist for any arbitrary interaction. We need to impose locality conditions.

Definition 1.8. *An interaction Φ is called:*

- ***Finite-range** if there is a radius $R > 0$ such that $\Phi(X) = 0$ whenever the diameter of the set X is greater than R.*
- ***Uniformly bounded** if there is a constant $M > 0$ such that $\|\Phi(X)\| \leq M$ for all finite sets X.*

Theorem 1.9 (Lieb, Robinson)**.** *If the interaction Φ is finite-range and uniformly bounded, then for any quasi-local observable $A \in \mathcal{A}_\Gamma$, the thermodynamic limit of the dynamics exists:*

$$\tau_t(A) = \lim_{\Lambda \uparrow \Gamma} \tau_t^\Lambda(A)$$

*The limit exists in the norm topology. This limiting evolution $\{\tau_t\}_{t \in \mathbb{R}}$ forms a **strongly continuous one-parameter group of *-automorphisms** of the quasi-local algebra $\mathcal{A}_\Gamma$. This means that for any $A \in \mathcal{A}_\Gamma$, the map $t \mapsto \tau_t(A)$ is continuous. The pair $(\mathcal{A}_\Gamma, \tau_t)$ is called a **C*-dynamical system**.*

The Generator of Dynamics (Derivation)

Definition 1.10 (Generator of Dynamics)**.** *From the general theory of strongly continuous one-parameter groups, there exists a* ***generator****,* δ*. This generator is a densely-defined, generally unbounded, closed operator. Its domain,* $\mathrm{Dom}(\delta)$*, consists of all observables* $A \in \mathcal{A}_\Gamma$ *for which the map* $t \mapsto \tau_t(A)$ *is differentiable. For any* $A \in \mathrm{Dom}(\delta)$*, the generator is defined as:*

$$\delta(A) = \frac{1}{i} \frac{d}{dt} \tau_t(A) \Big|_{t=0}$$

This generator δ *is a* ****-derivation****, meaning it satisfies the Leibniz rule:*

$$\delta(A^*) = \delta(A)^* \quad \textit{and} \quad \delta(AB) = \delta(A)B + A\delta(B)$$

In a finite system, this generator is simply the commutator with the Hamiltonian: $\delta(A) = [H_\Lambda, A]$*.*

1.1.4 States and Ground States

States on Quantum Systems

In a finite quantum system on Λ, a **state** is described by a **density matrix** ρ, which is a positive semi-definite matrix ($\rho \geq 0$) with unit trace ($\mathrm{Tr}(\rho) = 1$). The expectation value of an observable $A \in \mathcal{A}_\Lambda$ is then given by $\mathrm{Tr}(\rho A)$. A state is **pure** if ρ is a rank-one projection, $\rho = |\psi\rangle\langle\psi|$. This concept generalizes to the infinite C*-algebra $\mathcal{A}_\Gamma$.

Definition 1.11. *A* ***state*** *on the quasi-local algebra* $\mathcal{A}_\Gamma$ *is a linear functional* $\omega : \mathcal{A}_\Gamma \to \mathbb{C}$ *that is:*

1. ***Positive:*** $\omega(A^*A) \geq 0$ *for all* $A \in \mathcal{A}_\Gamma$*.*
2. ***Normalized:*** $\omega(\mathbf{1}) = 1$*.*

Ground States

At very low temperatures, physical systems tend to settle into their **ground state**, or the state of minimum energy. For a finite system with Hamiltonian H_Λ, the ground states are those whose density matrices ρ have support only on the eigenspace corresponding to the lowest eigenvalue of H_Λ. For the infinite system, we no longer have a well-defined Hamiltonian operator in the algebra. However, we have its generator of dynamics, δ.

Definition 1.12. *A state* ω *on the C*-dynamical system* $(\mathcal{A}_\Gamma, \tau_t)$ *is called a* ***ground state*** *if:*

$$\omega(A^*\delta(A)) \geq 0$$

for all $A \in \mathrm{Dom}(\delta)$*.*

This is a robust definition. For a finite system, this condition is equivalent to the standard definition of a ground state. Furthermore, if you take a sequence of ground states of finite systems H_{Λ_n} and take a weak-* limit as $\Lambda_n \to \Gamma$, the resulting limit state will be a ground state of the infinite system according to this definition.

1.1.5 Key Examples

We will keep two famous models in mind as we develop the theory.

1. **The AKLT Chain:** Introduced by Affleck, Kennedy, Lieb, and Tasaki (1987).
   - **Lattice:** A 1D chain, $\Gamma = \mathbb{Z}$.
   - **Spins:** A spin-1 system, meaning the local dimension is $d_x = 3$ for all $x \in \mathbb{Z}$.
2. **The Toric Code:** Introduced by Kitaev ($\sim$2003).
   - **Lattice:** The edges of the 2D square lattice, $\Gamma = E(\mathbb{Z}^2)$.
   - **Spins:** A qubit system, meaning the local dimension is $d_x = 2$ for all $x \in \Gamma$.

Both of these models have Hamiltonians built from finite-range, bounded interactions, and their properties, particularly their ground states, can be analyzed within the framework we have just established.

1.2 Lecture 2

1.2.1 Recap: Ground States of Infinite Systems

We recall the definition of a ground state for a C*-dynamical system $(\mathcal{A}_\Gamma, \tau_t)$ with generator δ.

Definition 1.13 (Ground State)**.** *A state ω on the quasi-local algebra $\mathcal{A}_\Gamma$ is a **ground state** for the dynamics generated by δ if*

$$\omega(A^*\delta(A)) \geq 0 \quad \forall A \in Dom(\delta).$$

It is a well-known result for finite-range interactions that the algebra of local observables, $\mathcal{A}_{local}$, forms a core for the generator δ. Therefore, it is sufficient to check this condition for all $A \in \mathcal{A}_{local}$.

This definition is equivalent to the standard physical definition for finite systems. A powerful way to motivate and understand this definition is through the Gelfand-Naimark-Segal (GNS) construction.

1.2.2 The GNS Representation

The GNS construction shows that any state on a C*-algebra can be represented as a vector state in some Hilbert space representation.

Theorem 1.14 (GNS Construction)**.** *Let $\mathcal{A}$ be any C*-algebra and let ω be a state on $\mathcal{A}$. Then, there exists a Hilbert space $\mathcal{H}_\omega$, a representation $\pi_\omega : \mathcal{A} \to B(\mathcal{H}_\omega)$, and a unit vector $\Omega_\omega \in \mathcal{H}_\omega$ (called the cyclic vector) with the following properties:*

1. *The state ω is recovered as a vector state:*
$$\omega(A) = \langle \Omega_\omega, \pi_\omega(A)\Omega_\omega \rangle \quad \forall A \in \mathcal{A}.$$
2. *The set of vectors $\pi_\omega(\mathcal{A})\Omega_\omega = \{\pi_\omega(A)\Omega_\omega \mid A \in \mathcal{A}\}$ is dense in $\mathcal{H}_\omega$.*

Moreover, this representation is unique up to unitary equivalence. That is, if $(\mathcal{H}_1, \pi_1, \Omega_1)$ and $(\mathcal{H}_2, \pi_2, \Omega_2)$ are two such representations for the same state ω, then there exists a unitary operator $U : \mathcal{H}_1 \to \mathcal{H}_2$ such that $U\Omega_1 = \Omega_2$ and $U\pi_1(A) = \pi_2(A)U$ for all $A \in \mathcal{A}$.

For finite-dimensional matrix algebras, all irreducible representations are unitarily equivalent. For infinite systems, this is not the case; there can be many unitarily inequivalent representations, and this is directly related to the classification of different physical phases of matter.

Symmetries in the GNS Representation

The GNS construction is particularly useful for understanding symmetries.

Proposition 1.15. *Let ω be a state on a C*-algebra $\mathcal{A}$, and let $\alpha : \mathcal{A} \to \mathcal{A}$ be a *-automorphism that leaves the state invariant, i.e., $\omega \circ \alpha = \omega$. Then α is implementable by a unitary operator in the GNS representation of ω.*

Proof Sketch. Let $(\mathcal{H}_\omega, \pi_\omega, \Omega_\omega)$ be the GNS representation for ω. We can construct a second representation $(\mathcal{H}_\omega, \pi_\omega \circ \alpha, \Omega_\omega)$. We check that this new triple also represents the state ω:

$$\langle \Omega_\omega, (\pi_\omega \circ \alpha)(A)\Omega_\omega \rangle = \omega(\alpha(A)) = \omega(A).$$

Since this new representation represents the same state ω on the same Hilbert space, by the uniqueness part of the GNS theorem, it must be unitarily equivalent to the original one. This means there exists a unitary operator U on $\mathcal{H}_\omega$ such that $\pi_\omega(\alpha(A)) = U^* \pi_\omega(A) U$ and $U\Omega_\omega = \Omega_\omega$. □

Dynamics and the GNS Hamiltonian

We can apply this result to the time evolution τ_t. A stationary state is one that is invariant under the dynamics, $\omega \circ \tau_t = \omega$. A simple calculation shows that any ground state is stationary:

$$\omega(\delta(A)) = \frac{1}{i}\frac{d}{dt}\omega(\tau_t(A))\Big|_{t=0} = 0.$$

Since a ground state ω is invariant under τ_t, the proposition above implies that the dynamics is implemented by a unitary group $\{U_t\}_{t\in\mathbb{R}}$ in the GNS Hilbert space.

$$\pi_\omega(\tau_t(A)) = U_t^* \pi_\omega(A) U_t \quad \text{and} \quad U_t\Omega_\omega = \Omega_\omega.$$

The strong continuity of the automorphism group τ_t implies that the unitary group U_t is also strongly continuous. By Stone's Theorem, such a group has a densely defined, self-adjoint generator H_ω, which we call the **GNS Hamiltonian**.

$$U_t = e^{-itH_\omega}.$$

1.2.3 Gapped Ground States

Proposition 1.16. *For a ground state ω, the GNS Hamiltonian H_ω is positive semi-definite, i.e., $H_\omega \geq 0$.*

Proof Sketch. The ground state condition $\omega(A^*\delta(A)) \geq 0$ can be rewritten in the GNS representation. For any $A \in \mathrm{Dom}(\delta)$:

$$0 \leq \omega(A^*\delta(A)) = \langle \Omega_\omega, \pi_\omega(A^*)\pi_\omega(\delta(A))\Omega_\omega \rangle = \langle \pi_\omega(A)\Omega_\omega, H_\omega \pi_\omega(A)\Omega_\omega \rangle.$$

Since the vectors $\pi_\omega(A)\Omega_\omega$ are dense in $\mathcal{H}_\omega$, this implies $H_\omega \geq 0$. The spectrum of the GNS Hamiltonian is therefore contained in $[0, \infty)$, and the ground state vector Ω_ω is an eigenvector with eigenvalue 0. □

This motivates the definition of a spectral gap.

Definition 1.17 (Gapped Ground State)**.** *A ground state ω is said to be **gapped** if there exists a constant $\gamma > 0$ such that the spectrum of the GNS Hamiltonian H_ω is contained in $\{0\} \cup [\gamma, \infty)$.*

This physical condition on the energy spectrum has an equivalent algebraic formulation.

Proposition 1.18. *A ground state ω has a spectral gap of at least $\gamma > 0$ if and only if*

$$\omega(A^*\delta(A)) \geq \gamma\,\omega(A^*A)$$

for all $A \in Dom(\delta)$ such that $\omega(A) = 0$. The condition $\omega(A) = 0$ is equivalent to the GNS vector $\pi_\omega(A)\Omega_\omega$ being orthogonal to the ground state vector Ω_ω.

1.2.4 Pure and Mixed States

Definition 1.19. *A state ω on a C*-algebra $\mathcal{A}$ is **pure** if it cannot be written as a non-trivial convex combination of other states. That is, if $\omega = t\omega_1 + (1-t)\omega_2$ for $t \in (0,1)$ implies that $\omega_1 = \omega_2 = \omega$. If a state is not pure, it is called **mixed**.*

A state is pure if and only if its GNS representation is irreducible. The set of all ground states for a given interaction forms a convex set. The physically relevant ground states are often the extreme points of this set, which are the pure ground states.

1.2.5 Examples

The AKLT Model

The Affleck-Kennedy-Lieb-Tasaki (AKLT) model is a spin-1 chain on $\Gamma = \mathbb{Z}$. The interaction is nearest-neighbor and given by the projection onto the total spin-2 subspace of two adjacent spins.

$$\Phi(\{x, x+1\}) = P^{(2)}_{x,x+1} = \frac{1}{3}\mathbb{I} + \frac{1}{2}\boldsymbol{S}_x \cdot \boldsymbol{S}_{x+1} + \frac{1}{6}(\boldsymbol{S}_x \cdot \boldsymbol{S}_{x+1})^2$$

where $\boldsymbol{S}_x = (S^x_x, S^y_x, S^z_x)$ are the spin-1 matrices at site x.

- There exists a **unique** ground state ω for this model.
- This ground state is "frustration-free," meaning it has zero energy with respect to every local term in the Hamiltonian: $\omega(P^{(2)}_{x,x+1}) = 0$ for all x.
- The ground state is proven to be gapped, with a gap known to be at least ≈ 0.1.

The Toric Code

The Toric Code model is defined on the edges of the 2D square lattice $\mathbb{Z}^2$, with a qubit ($d_x = 2$) on each edge.

- The set of ground states is not unique. Instead, it consists of 4 unitarily inequivalent classes, corresponding to different "superselection sectors."
- These sectors correspond to the vacuum state (no charge), an electric charge sector, a magnetic charge sector, and a fused dyonic charge sector.

- Each sector contains a ground state, and these states cannot be distinguished by any local observable.

1.2.6 Outlook: Local Approximations

A key tool for analyzing locality is the conditional expectation. Given the tracial state on the whole lattice, $\rho = \bigotimes_{x\in\Gamma} \frac{1}{d_x}\mathbf{1}_x$, we can define a map $\pi_\Lambda : \mathcal{A}_\Gamma \to \mathcal{A}_\Lambda$ for any finite region $\Lambda \subset \Gamma$. This map acts as the identity on $\mathcal{A}_\Lambda$ and takes the partial trace over the complement $\Gamma \setminus \Lambda$. For any $A \in \mathcal{A}_\Gamma$, we have

$$\lim_{\Lambda\uparrow\Gamma} \|A - \pi_\Lambda(A)\| = 0.$$

This provides a canonical way to approximate any quasi-local observable with a sequence of strictly local ones, which will be a crucial tool going forward.

1.3 Lecture 3

1.3.1 Quantifying Locality

Conditional Expectation and Commutators

We begin with a recap of the conditional expectation map $\pi_\Lambda : \mathcal{A}_\Gamma \to \mathcal{A}_\Lambda \subset \mathcal{A}_\Gamma$, which is defined by taking the partial trace over the complement of Λ. For any $A \in \mathcal{A}_\Gamma$ and any increasing sequence of finite volumes $\Lambda_n \uparrow \Gamma$, we have $\lim_n \|\pi_{\Lambda_n}(A) - A\| = 0$.

The error in this approximation can be related to commutators. For a finite-dimensional algebra, the normalized trace can be represented by averaging over the unitary group: $(\frac{1}{D}\operatorname{Tr} A)\mathbf{1} = \int UAU^* dU$. A similar representation for the partial trace gives:

$$A - \pi_\Lambda(A) = \int_{U(\mathcal{H}_{\Gamma\setminus\Lambda})} dU\, U\,[U^*, A]\,,$$

where the integral is over the unitary group acting on the complement of Λ. This leads to a useful bound.

Proposition 1.20. *The norm of the error of the local approximation is bounded by the norm of commutators with observables in the complement:*

$$\|A - \pi_\Lambda(A)\| \leq \sup_{\substack{B \in \mathcal{A}_{\Gamma\setminus\Lambda} \\ \|B\|=1}} \|[A, B]\| \leq 2\|A - \pi_\Lambda(A)\|.$$

Algebras of Localized Observables We can use this connection to define new Banach algebras of observables with specific decay properties.

Definition 1.21 (Decay Function)**.** *A **decay function** is a function $g : [0, \infty) \to (0, \infty)$ that is non-increasing and vanishes at infinity, $\lim_{r\to\infty} g(r) = 0$.*

Given a decay function g and an increasing sequence of finite volumes (Λ_n), we can define a new norm.

Definition 1.22 (g-local Norm)**.** *The **g-local norm** of an observable A with respect to a sequence (Λ_n) is*

$$\|A\|_{(\Lambda_n),g} = \|A\| + \sup_n \frac{\|A - \pi_{\Lambda_n}(A)\|}{g(n)}.$$

The completion of the local observables $\mathcal{A}_{local}$ with respect to this norm, denoted $\mathcal{A}_{(\Lambda_n),g}$, is a Banach algebra.

If we choose Λ_n to be balls of increasing radius, $\Lambda_n = b_x(n)$, we can define an algebra $\mathcal{A}^g_\Gamma = \bigcup_{x\in\Gamma} \mathcal{A}^{(b_x(n)),g}$. If g satisfies a subadditivity condition like $g(n)g(m) \leq Cg(n + m)$, then the norms $\|\cdot\|_{x,g}$ are equivalent for all $x, y \in \Gamma$, and the union is trivial. We call $\mathcal{A}^g_\Gamma$ the algebra of **g-local observables**.

Definition 1.23 (Almost Local Observables)**.** *For the family of power-law decay functions $g_\alpha(r) = (1+r)^{-\alpha}$ for $\alpha > 0$, the algebra of* ***almost local observables*** *is defined as the intersection of all such g-local algebras:*

$$\mathcal{A}_\Gamma^{al} = \bigcap_{\alpha>0} \mathcal{A}_\Gamma^{g_\alpha}.$$

This is a Fréchet space whose topology is defined by the family of norms $\{\|\cdot\|_{g_\alpha}\}_{\alpha>0}$.

1.3.2 Banach Spaces of Interactions

We can similarly define normed spaces for interactions. We first need a suitable class of decay functions.

Definition 1.24 (F-function)**.** *A decay function g is called an* ***F-function*** *if it satisfies:*

1. ***Uniform Summability:*** *The function is summable over the lattice, uniformly in the starting point:*

$$\|g\|_1 = \sup_{x\in\Gamma} \sum_{y\in\Gamma} g(d(x,y)) < \infty.$$

2. ***Convolution Inequality:*** *There exists a constant $C_g > 0$ such that*

$$\sum_{z\in\Gamma} g(d(x,z))g(d(z,y)) \leq C_g g(d(x,y)), \quad \forall x, y \in \Gamma.$$

A common example of F-functions are of the form $g(r) = (1+r)^{-\xi}e^{-ar}$ for sufficiently large ξ.

Definition 1.25 (Interaction Norm)**.** *For an interaction Φ and an F-function g, we define the norm:*

$$\|\Phi\|_g = \sup_{x,y\in\Gamma} \frac{1}{g(d(x,y))} \sum_{\substack{X\in P_0(\Gamma)\\ x,y\in X}} \|\Phi(X)\|.$$

The space of all interactions for which this norm is finite is a Banach space, which we denote by $\mathcal{B}_g$. If the interaction is time-dependent, we can define spaces like $\mathcal{B}_g([0,1])$ by taking the supremum of $\|\Phi_t\|_g$ over the time interval.

1.3.3 Lieb-Robinson Bounds

The dynamics of systems with suitably decaying interactions are constrained by Lieb-Robinson (LR) bounds, which establish an effective "speed of light" for the propagation of information.

Dynamics for Time-Dependent Hamiltonians For a finite system with a time-dependent Hamiltonian $H(s)$, the dynamics is governed by the Schrödinger equation $\frac{d}{ds}U(s,0) = -iH(s)U(s,0)$. In the Heisenberg picture, the time evolution of an observable A is given by $\tau_{s,0}(A) = U^*(s,0)AU(s,0)$, which satisfies $\frac{d}{ds}\tau_{s,0}(A) = i\tau_{s,0}([H(s), A])$.

The Main Bound

Theorem 1.26 (Basic Lieb-Robinson Bound)**.** *Let $\Phi \in \mathcal{B}_g(I)$ be a time-dependent interaction where g is an F-function. For any two disjoint finite sets $X, Y \subset \Gamma$ ($X \cap Y = \emptyset$), any $A \in \mathcal{A}_X$, $B \in \mathcal{A}_Y$, and any $s,t \in I$, the following bound holds:*

$$\|[\tau^{\Phi}_{t,s}(A), B]\| \leq 2\|A\|\|B\| \left(e^{2\left|\int_s^t \|\Phi_r\|_g dr\right|} - 1\right) D(X,Y),$$

where $D(X,Y) = \frac{1}{C_g}\sum_{x\in X}\sum_{y\in Y} g(d(x,y))$.

This theorem shows that the commutator between a time-evolved observable and a distant observable is small, with the decay governed by the function g.

Consequences and the Light Cone

- **Time-Independent Case:** If Φ is time-independent, the time-dependent factor becomes $(e^{2|t-s|\|\Phi\|_g} - 1)$.
- **Exponential Decay:** If $g(r)$ decays exponentially, e.g., $g(r) \sim e^{-ar}$, then the spatial factor $D(X,Y)$ also decays exponentially with the distance between the sets: $D(X,Y) \leq C\min(|X|,|Y|)e^{-ad(X,Y)}$. This leads to the famous "light cone" picture, where the support of a time-evolved local observable $\tau_t(A)$ is essentially contained within a region that grows linearly with time. More precisely, the part of $\tau_t(A)$ outside a ball of radius $v|t|$ (where $v \sim \|\Phi\|_g/a$) is exponentially small.

Existence of Infinite-Volume Dynamics LR bounds are the key to proving the existence of the dynamics in the thermodynamic limit.

Theorem 1.27. *Let $X \subset \Lambda_0 \subset \Lambda_1$ be finite volumes. For any $A \in \mathcal{A}_X$, the difference between the dynamics in the two volumes is bounded:*

$$\|\tau_t^{\Lambda_1}(A) - \tau_t^{\Lambda_0}(A)\| \leq C\|A\|\|\Phi\|_g e^{2\|\Phi\|_g|t|} \sum_{x\in X}\sum_{y\in\Lambda_1\setminus\Lambda_0} g(d(x,y)).$$

As $\Lambda_0, \Lambda_1 \to \Gamma$, the sum on the right-hand side goes to zero because of the summability of g. This shows that for any local observable A, the sequence $\{\tau_t^{\Lambda_n}(A)\}_{n\in\mathbb{N}}$ is a Cauchy sequence.

Corollary 1.28. *For any interaction $\Phi \in \mathcal{B}_g$, the thermodynamic limit $\lim_n \tau_t^{\Lambda_n}(A) = \tau_t(A)$ exists for all $A \in \mathcal{A}_{local}$. This limit extends to a strongly continuous group of automorphisms on the entire quasi-local algebra $\mathcal{A}_\Gamma$.*

1.3.4 Introduction to Quasi-Adiabatic Evolution

A powerful application of LR bounds is the construction of the quasi-adiabatic evolution, introduced by Hastings. This is a tool for relating systems with different Hamiltonians. Consider a smooth path of interactions Φ_s for $s \in [0,1]$. We can define a special norm on this path:

$$|||\Phi|||_1 = \sup_{s\in[0,1]} \sup_{x,y\in\Gamma} \frac{1}{g(d(x,y))} \sum_{\substack{X\in P_0(\Gamma)\\ x,y\in X}} (\|\Phi_s(X)\| + |X|\|\dot{\Phi}_s(X)\|).$$

The quasi-adiabatic evolution is generated by an effective interaction, often called the Hastings generator, which is constructed from the path derivative $\frac{d}{ds}\Phi_s$. Its terms are given by an integral over time-evolved derivatives:

$$\mathcal{D}_X^{(\Phi)}(s) = \int_{-\infty}^{\infty} W(t)\tau_t^{\Phi_s}\left(\frac{d}{ds}\Phi_s(X)\right) dt,$$

where $W(t)$ is a suitable rapidly decaying weight function. This construction allows one to build an automorphism that connects the ground states of Hamiltonians at different points along the path, which has profound implications for the classification of gapped phases of matter.

1.4 Lecture 4

1.4.1 Defining Gapped Phases of Matter

The central goal is to develop a rigorous framework for classifying phases of matter in quantum spin systems, particularly those characterized by a spectral gap above the ground state. This leads to the notion of a "gapped phase." We will explore two primary, and ultimately equivalent, definitions for when two Hamiltonians (or their ground states) belong to the same phase.

Let $\mathfrak{A}$ be the C*-algebra of quasi-local observables on a discrete metric space (Γ, d). We consider interactions Φ which are collections of local Hamiltonian terms, for which we can define suitable norms that capture spatial decay properties. A key example is the space $B_{a,\theta}$ of interactions with exponential decay.

Suppose we have two interactions, Φ_0 and Φ_1, belonging to such a class, e.g., $\Phi_0, \Phi_1 \in B_{a,\theta}$ for some $a, \theta > 0$. Assume these interactions give rise to Hamiltonians with unique, gapped ground states, denoted ω_0 and ω_1, respectively. The fundamental question is: when should we consider these systems to be in the same phase?

Definition 1.29 (Equivalence via Gapped Path (Chen-Gu-Wen '10))**.** *Two interactions Φ_0 and Φ_1 are in the same gapped phase, written $\Phi_0 \sim \Phi_1$, if there exists a C^1-curve of interactions $\Phi : [0,1] \to B_{a',\theta'}$ for some $a', \theta' > 0$, such that:*

1. *The curve interpolates between the two interactions: $\Phi(0) = \Phi_0$ and $\Phi(1) = \Phi_1$.*
2. *The spectral gap of the Hamiltonian $H(\Phi(s))$ is uniformly bounded from below along the entire path. That is, there exists a $\gamma > 0$ such that for all $s \in [0,1]$, the gap of $H(\Phi(s))$ is at least γ.*

This defines an equivalence relation on the set of gapped interactions. The differentiability condition can be relaxed, but it is mathematically convenient for proving the equivalence to the next definition.

An alternative perspective emphasizes the ground states themselves, rather than the Hamiltonians that produce them. This leads to a definition based on the existence of a special type of automorphism of the observable algebra.

Definition 1.30 (Equivalence via Automorphism)**.** *Two ground states ω_0 and ω_1 (or sets of ground states) are in the same phase if there exists a continuous path of automorphisms $\alpha_s : \mathfrak{A} \to \mathfrak{A}$ for $s \in [0,1]$, generated by a suitable parameter-dependent, quasi-local interaction $\Psi(s)$, such that $\omega_1 = \omega_0 \circ \alpha_1$. The automorphism path α_s is the solution to the differential equation for the dynamics, where s plays the role of time and $\Psi(s)$ is the generator.*

1.4.2 The Quasi-Adiabatic Evolution

The bridge between the two definitions is the **quasi-adiabatic evolution**, a technique pioneered by M. Hastings. It provides a method to construct an automorphism that "follows" the ground state of a system as its Hamiltonian is slowly changed, provided the spectral gap never closes.

Constructing the Hastings Generator Given a C^1-curve of interactions $\Phi(s)$, we construct a new, time-dependent interaction, $\Psi(s)$, known as the **Hastings generator**. First, we define an intermediate object $\tilde{\Psi}(s) \in \mathfrak{A}$:

$$\tilde{\Psi}(s) = i \int_{-\infty}^{\infty} W_a(t)\, \tau_t^{\Phi(s)} \left(\frac{d\Phi(s)}{ds} \right) dt$$

where:

- $\tau_t^{\Phi(s)}$ is the Heisenberg dynamics generated by the (time-independent) Hamiltonian $H(\Phi(s))$.
- $d\Phi(s)/ds$ is the derivative of the interaction path.
- $W_a(t)$ is a rapidly decaying, odd weight function in $L^1(\mathbb{R})$. It is constructed to satisfy certain moment conditions and has exponential decay $|W_a(t)| \leq Ce^{-\alpha|t|}$. The parameter a controls the decay rate and must be chosen sufficiently small relative to the uniform spectral gap γ.

This $\tilde{\Psi}(s)$ is a quasi-local observable, but it is not an interaction in the sense of being a sum of strictly local terms. To obtain a proper interaction, we decompose it using conditional expectations. This process effectively "cuts up" the quasi-local observable $\tilde{\Psi}(s)$ into a sum of local terms, $\Psi(s) = \sum_X \Psi_X(s)$, with rapidly decaying norms, yielding a well-defined interaction. The local terms can be defined via a telescoping sum:

$$\begin{aligned} \Psi_{X_1}(s) &= E_{X_1}(\tilde{\Psi}(s)) \\ \Psi_{X_n}(s) &= E_{X_n}(\tilde{\Psi}(s)) - E_{X_{n-1}}(\tilde{\Psi}(s)) \quad \text{for } n > 1 \end{aligned}$$

where X_n is an increasing sequence of finite regions (e.g., balls of increasing radius) and E_{X_n} is the conditional expectation onto the local algebra $\mathfrak{A}_{X_n}$. Summing these terms recovers the original $\tilde{\Psi}(s)$ in the limit.

1.4.3 Equivalence of Definitions and Stability of the Gap

The Hastings generator is precisely the object needed to connect the two definitions of a phase.

Theorem 1.31 (Equivalence of Definitions)**.** *Suppose two interactions Φ_0 and Φ_1 are equivalent according to Definition 1, connected by a C^1-path $\Phi(s)$ with a uniform spectral gap $\gamma > 0$. Let $\Psi(s)$ be the Hastings generator constructed from this path with a suitable choice of weight function $W_a(t)$ (where a is small*

enough with respect to γ). Then the family of automorphisms α_s, generated by solving the dynamics from time 0 to s with generator $\Psi(s)$, satisfies

$$\omega_s = \omega_0 \circ \alpha_s \quad \textit{for all } s \in [0,1],$$

where ω_s is the unique gapped ground state of $H(\Phi(s))$. In particular, $\omega_1 = \omega_0 \circ \alpha_1$, establishing equivalence under Definition 2.

Remark 1.32. *This construction preserves any symmetries of the original path $\Phi(s)$. If each $\Phi(s)$ is invariant under a certain symmetry group, the resulting generator $\Psi(s)$ and automorphism α_s will also be invariant under that symmetry. This is crucial for the classification of Symmetry-Protected Topological (SPT) phases.*

Stability of the Spectral Gap A fundamental question is whether the existence of a spectral gap is a stable property. The following theorem addresses this for a broad class of frustration-free systems.

Assumptions for Gap Stability Let $H_0 = \sum_x h_x$ be the unperturbed, frustration-free Hamiltonian.

1. **Lattice Regularity:** The underlying metric space (Γ, d) is of 'new-regularity' and admits 'decomposable partitions', which prevents points from bunching up and allows for certain coloring schemes. The lattice $\mathbb{Z}^\nu$ is a standard example.
2. **Finite Range:** The unperturbed interaction is finite-range. There exists $R > 0$ such that $h_x \in \mathfrak{A}_{B(x,R)}$ for all $x \in \Gamma$.
3. **Frustration-Freeness:** There exists a ground state ω_0 such that $\omega_0(h_x) = 0$ for all $x \in \Gamma$ (since $h_x \geq 0$).
4. **Uniform Boundedness:** The local terms are uniformly bounded: $\sup_x \|h_x\| < \infty$.
5. **Bulk Gap:** The GNS Hamiltonian for the infinite-volume ground state ω_0 has a spectral gap $\gamma_0 > 0$. $\mathrm{Spec}(H_{\omega_0}) \subset \{0\} \cup [\gamma_0, \infty)$.
6. **Vanishing of Edge Modes:** The gap of the finite-volume Hamiltonian $H_{0,\Lambda} = \sum_{x \in \Lambda} h_x$ may close as $|\Lambda| \to \infty$, but not too quickly (e.g., polynomially). This allows for gapless edge modes in the thermodynamic limit of half-spaces, while the bulk remains gapped.
7. **Perturbation Decay:** The perturbation Φ has summable, exponentially decaying tails, i.e., it belongs to a space like $B_{a,\theta}$.
8. **Local Topological Quantum Order (LTQO):** The ground state subspace of finite-volume Hamiltonians is locally indistinguishable from the global ground state. For any local observable A supported in a ball $B(x,k)$, its matrix elements are constrained. For any finite volume Λ containing $B(x,m)$ with $m > k$:

$$\|\Pi_\Lambda A \Pi_\Lambda - \omega_0(A)\Pi_\Lambda\| \leq \|A\| \cdot G(m-k)$$

where Π_Λ is the projection onto the ground state space of $H_{0,\Lambda}$, and the function G decays fast enough to have a finite moment of a sufficiently high order.

Theorem 1.33 (Stability of the Bulk Spectral Gap)**.** *Assume the unperturbed model H_0 and the perturbation Φ satisfy assumptions 1-8. Then the spectral gap is stable. For any target gap $\gamma \in (0, \gamma_0)$, there exists a constant $\beta > 0$ such that for all perturbation strengths $|s| \leq \frac{\gamma_0 - \gamma}{\beta}$, the perturbed Hamiltonian $H(s) = H_0 + s\Phi$ has a gapped ground state ω_s with a spectral gap of at least γ. Furthermore, the state ω_s is connected to ω_0 by the quasi-adiabatic evolution.*

1.4.4 Application: O(n) Spin Chains

The theoretical framework can be applied to concrete physical models. A rich example is the family of $O(n)$-invariant spin-1 chains, which generalize the AKLT model. The Hamiltonian for the original AKLT model ($n = 3$) involves a projector $P^{(2)}$ onto the total spin-2 subspace for two neighboring spin-1's. This interaction can be written in an equivalent form using the swap operator $T_{x,x+1}$ and a rank-1 projector $Q_{x,x+1}$:

$$P^{(2)}_{x,x+1} = \frac{1}{3}\mathbb{1} + \frac{1}{2}\mathbf{S}_x \cdot \mathbf{S}_{x+1} + \frac{1}{6}(\mathbf{S}_x \cdot \mathbf{S}_{x+1})^2 \cong \frac{1}{2}(\mathbb{1} + T_{x,x+1} - 2Q_{x,x+1})$$

This form motivates a general family of $O(n)$-invariant nearest-neighbor interactions for n-dimensional spins:

$$H = \sum_x (uT_{x,x+1} + vQ_{x,x+1})$$

where $u, v \in \mathbb{R}$ and Q is the projection onto the maximally entangled state $\psi = \frac{1}{\sqrt{n}} \sum_{\alpha=1}^{n} |\alpha, \alpha\rangle$ in some orthonormal basis $\{|\alpha\rangle\}$.

This family of models exhibits a rich phase structure. Key points and regions in the phase diagram include:

- A Bethe ansatz integrable point at $v = -2nu/(n-2)$ for $n \geq 3$ (Reshetikhin, 1983).
- A frustration-free point at $v = -2u$. This point is gapped for all n. The stability theorem ensures the gap persists in an open neighborhood, defining the stable, red "Haldane-like" phase. The ground state is unique for odd n but doubly degenerate (2-periodic) for even n (Nachtergaele-Sims-Young, 2022).
- At $u = 0, v = -1$, the model is equivalent to the $SU(n)$ Temperley-Lieb chain, which is dimerized and gapped for all $n \geq 3$ (Aizenman, et al. 2020).

The framework of gapped phases allows for a sharp distinction between these regions. For example, in the Haldane-like phase ($v \approx -2u$):

- For **odd n** (like the AKLT model), the system has a unique gapped ground state in a non-trivial SPT phase protected by symmetries like inversion and $SO(n)$.
- For **even n**, the ground states spontaneously break the full $O(n)$ symmetry down to $SO(n)$ and also break translation symmetry.

The yellow dimerized phase, in contrast, preserves the full $O(n)$ symmetry while breaking translation symmetry. Because the red (for even n) and yellow phases have different symmetry breaking patterns, they cannot be connected by a gapped path that preserves the $O(n)$ symmetry. Therefore, they represent distinct phases of matter, and a phase transition must occur between them. This illustrates the power of the quasi-adiabatic evolution framework in rigorously classifying phases of matter.

2 Mike Hopkins: Lattice Models and TQFTs

Contents

2.1 Lecture 1

2.1.1 Gapped Hamiltonian Systems

Consider a metric space analogous to $\mathbb{R}^n$, and let $\ell \in L$ index a family of Hilbert spaces $\mathcal{H}_\ell$, each of constant finite dimension. Let H denote a Hamiltonian acting on the tensor product space

$$\bigotimes_{\ell \in L} \mathcal{H}_\ell.$$

We are interested in the class of Hamiltonians that possess a spectral gap. Define $\mathcal{L}(X)$ to be the collection of such gapped Hamiltonian systems supported on a topological space X.

We can define an operation

$$\mathcal{L}(X) \times \mathcal{L}(X) \longrightarrow \mathcal{L}(X), \quad (H_1, H_2) \mapsto H_1 \otimes 1 + 1 \otimes H_2,$$

which corresponds to taking the external sum of independent systems. The new Hamiltonian $H_1 \boxtimes H_2$ acts on the space

$$\left(\bigotimes_{\ell_1} \mathcal{H}^{(1)}_{\ell_1} \right) \otimes \left(\bigotimes_{\ell_2} \mathcal{H}^{(2)}_{\ell_2} \right).$$

Definition 2.1. *A Hamiltonian model $H \in \mathcal{L}(X)$ is called **invertible** if there exists $H' \in \mathcal{L}(X)$ such that*

$$H \otimes H' \simeq 1,$$

where 1 denotes the trivial system defined by $\mathcal{H}_\ell = \mathbb{C}$ and $H = 0$.

Let $\mathcal{K}(X) \subset \mathcal{L}(X)$ denote the subcollection of invertible systems.

Theorem 2.2 (Kitaev)**.** *The space $\mathcal{K}(\mathbb{R}^n)$ defines a spectrum:*

$$\Omega \mathcal{K}(\mathbb{R}^{n+1}) \simeq \mathcal{K}(\mathbb{R}^n).$$

Problem 2.3. *Identify the spectrum represented by $\mathcal{K}(\mathbb{R}^n)$.*

2.1.2 Spectra and Cohomology Theories

A *cohomology theory* is a contravariant functor $X \mapsto E(X)$ from topological spaces to abelian groups, satisfying exactness. In particular, for $X = U \cup V$,

the following Mayer-Vietoris exact sequence holds:

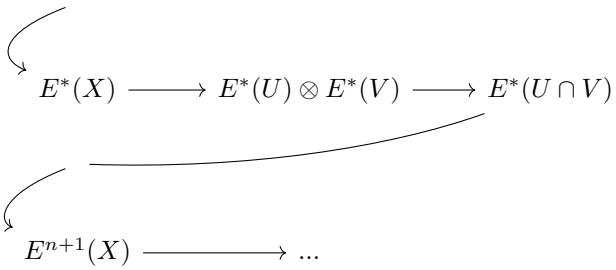

Theorem 2.4 (Brown Representability)**.** *Every cohomology theory E^* is representable by a spectrum $\{\underline{E}^n\}$, meaning*

$$E^n(X) \cong [X, \underline{E}^n], \quad \underline{E}^n \simeq \Omega \underline{E}^{n+1}.$$

Define the spectrum $\mathcal{K}$ by

$$\underline{\mathcal{K}}^n := \mathcal{K}(\mathbb{R}^{n-1}) \quad (n \geq 1), \qquad \underline{\mathcal{K}}^n \simeq \Omega \underline{\mathcal{K}}^{n+1}.$$

For $\mathbb{R}^0$, a point, we have:

$$\mathcal{L}(\mathbb{R}^0) \cong \bigsqcup_{p,q} \mathrm{Gr}(\mathbb{C}^{p|q}).$$

Each Hamiltonian acts on a graded vector space $\mathbb{C}^{p|q}$, and the invertible systems correspond to graded lines (superlines):

$$\mathcal{K}(\mathbb{R}^0) \cong \mathbb{Z}_2 \times \mathbb{C}P^0.$$

Hence,

$$\underline{\mathcal{K}}^1 \cong \mathbb{Z}_2 \times \mathbb{C}P^0, \qquad \mathcal{K}^0 = \Omega(\mathbb{Z}_2 \times \mathbb{C}P^0) \cong S^1 \simeq \mathbb{C}^\times.$$

Definition 2.5. *The **homotopy groups** of a spectrum E are defined as*

$$\pi_n(E) := \pi_n(\underline{E}^0) = [S^n, \underline{E}^0].$$

Definition 2.6 (Brown–Comenetz Dual)**.** *Given a spectrum E and an injective abelian group A, define the dual spectrum $I_A(E)$ via*

$$I_A(E)^0(X) := \mathrm{Hom}(E_0(X), A),$$

where

$$E_0(X) := \varinjlim_{k \to \infty} \pi_{n+k}(E_k \wedge X).$$

Let $A = \mathbb{C}^\times$. Then

$$I_{\mathbb{C}^\times}(E) = \mathrm{Maps}(E, I_{\mathbb{C}^\times}).$$

Example 2.7. *Let $E = S$, the sphere spectrum. Then:*

$$\begin{aligned} \pi_i I_{\mathbb{C}^\times} &= 0 \quad \textit{for } i > 0, \\ \pi_0 I_{\mathbb{C}^\times} &= \mathrm{Hom}(\pi_1 S^0, \mathbb{C}^\times) \cong \mathbb{Z}_2, \\ \pi_{-1} I_{\mathbb{C}^\times} &= \mathrm{Hom}(\pi_0 S^0, \mathbb{C}^\times) \cong \mathbb{C}^\times. \end{aligned}$$

More generally,

$$E_1(I_{\mathbb{C}^\times}) = \mathrm{Hom}(\pi_0 E, \mathbb{C}^\times).$$

There is a sequence of duals:

$$I_{\mathbb{Z}} \to I_{\mathbb{C}} \to I_{\mathbb{C}^\times} \to \Sigma I_{\mathbb{Z}}.$$

Conjecture 2.1 (Kitaev)**.**

1. *There exists a commuting diagram of spectra:*

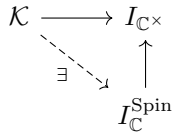

2. *This map is an equivalence of spectra.*

Conjecture 2.2. *The spectrum $\mathcal{K}$ is equivalent to the Anderson or Brown–Comenetz dual of M*Spin*.*

The Spin-invariance of the map $\mathcal{K} \to I_{\mathbb{C}^\times}$ corresponds physically to unitarity, or reflection positivity. It also implies emergent Lorentz symmetry in the associated quantum system.

2.2 Lecture 2

2.2.1 Introduction and Central Objects

Framing the Discussion This lecture series aims to explore the deep connections between the physical theory of gapped phases of matter and the mathematical framework of homotopy theory, particularly as motivated by the Kitaev conjecture. The goal is to take lessons from the homotopy theory side and formulate clear questions about the Hamiltonian lattice model side, and vice versa. Even for those unfamiliar with one side of this correspondence, the hope is that there is still something to be learned from the interplay of ideas. As the lectures progress, the explicit use of homotopy theory will lessen in favor of focusing on these core questions.

Defining the Spaces of Phases We begin by defining the primary objects of study, which represent the physical systems we wish to classify.

Definition 2.8. *Let (Γ, d) be a discrete metric space, such as $\mathbb{R}^n$ with a lattice structure.*

1. *We denote by $L(\Gamma)$ the space of all **gapped Hamiltonian lattice systems** on Γ.*
2. *We denote by $K(\Gamma) \subset L(\Gamma)$ the subspace of **invertible** gapped Hamiltonian systems. Invertibility is understood in the sense of stacking, where a system and its inverse can be combined to form a trivial system, up to some local equivalence.*

From these spaces, which encode the physics, we can construct a central object in algebraic topology: a **spectrum**. A spectrum is a sequence of spaces $\{E_n\}$, where each space is (up to homotopy) the loop space of the next, $E_n \simeq \Omega E_{n+1}$. Such an object represents a generalized (co)homology theory.

Definition 2.9. *The **Kitaev Spectrum**, denoted K, is the spectrum whose n-th space is given by the space of invertible phases on $\mathbb{R}^n$:*

$$K_n := K(\mathbb{R}^n)$$

Remark 2.10 (On Indexing)**.** *In previous discussions, the n-th space might have been denoted $K(\mathbb{R}^{n-1})$. The index is shifted here to conform with the notation used by Alexei Kitaev. While the previous indexing was convenient for certain homotopy-theoretic statements, this convention is more convenient for other purposes, including direct comparison with Kitaev's work. The path-components of the space K_n, denoted $\pi_0(K_n)$, are interpreted as the distinct phases of n-dimensional, invertible, gapped Hamiltonian systems (with no extra imposed symmetries).*

2.2.2 The Kitaev Conjecture

The central conjecture, proposed by Alexei Kitaev, provides a precise homotopy-theoretic classification of these phases of matter. It identifies the physically-defined Kitaev spectrum K with a well-known spectrum from algebraic topology related to spin bordism and Anderson duality.

Conjecture 2.3 (Kitaev)**.** *The Kitaev spectrum K is equivalent to the spectrum of maps from the spin bordism spectrum, $MSpin$, to a shifted Anderson dual of the integers. There are two closely related formulations for these fermionic systems:*

1. $K \simeq Maps(MSpin, \Sigma^2 I_{\mathbb{Z}})$
2. $K \simeq Maps(MSpin, \Sigma I_{\mathbb{C}^*})$

where Σ^k denotes the k-fold suspension, and $I_{\mathbb{Z}}$ and $I_{\mathbb{C}^*}$ are the Anderson duals to the Eilenberg-MacLane spectra for the integers and the circle group, respectively.

Remark 2.11 (Bosonic and Other Systems)**.** *The use of $MSpin$ specifies that we are classifying fermionic systems. For bosonic systems, one would replace $MSpin$ with MSO. For systems with other symmetries, such as time-reversal, one would use other bordism theories, like one of the Pin or MO theories.*

This conjecture implies that the classification of these physical systems can be computed using the tools of algebraic topology. The phases of n-dimensional systems, $\pi_0(K_n)$, are given by homotopy classes of maps:

$$\begin{aligned}\pi_0(K_n) &= \pi_0(\Sigma^n K) \simeq \pi_0(\mathrm{Maps}(MSpin, \Sigma^n \Sigma I_{\mathbb{C}^*})) \\ &= [S^0, \mathrm{Maps}(MSpin, \Sigma^{n+1} I_{\mathbb{C}^*})]_{\mathrm{hTop}} \\ &= [MSpin, \Sigma^{n+1} I_{\mathbb{C}^*}]_{\mathrm{hTop}} \\ &= \mathrm{Hom}(\pi_{n+1}(MSpin), \mathbb{C}^*)\end{aligned}$$

This remarkable statement asserts that the phases of n-dimensional invertible fermionic systems are classified by the $\mathbb{C}^*$-valued **cobordism invariants** of closed $(n+1)$-dimensional spin manifolds.

Spin Bordism Groups To understand the classification, we need the homotopy groups of the spectrum $MSpin$. These are the spin bordism groups of a point, $\pi_k(MSpin) = \Omega_k^{Spin}$, which are well-known and computed by Anderson, Brown, and Peterson in the 1960s.

Definition 2.12 (The MSpin Spectrum)**.** *$MSpin$ is a spectrum representing spin bordism. As a homology theory, for a space X, $MSpin_k(X)$ is the group of equivalence classes of maps $f : M^k \to X$ from k-dimensional closed spin manifolds, where two are equivalent if they form the boundary of a map from an $(k+1)$-dimensional spin manifold with boundary. As a cohomology theory, it classifies families of manifolds moving through cobordisms over a base space.*

The coefficient groups are:

- $\pi_0(MSpin) = \mathbb{Z}$ (Counting oriented points)
- $\pi_1(MSpin) = \mathbb{Z}_2$ (Generated by the circle with the non-bounding spin structure)
- $\pi_2(MSpin) = \mathbb{Z}_2$ (Generated by the torus with the non-bounding spin structure)
- $\pi_3(MSpin) = 0$
- $\pi_4(MSpin) = \mathbb{Z}$ (Given by the index of the Dirac operator, the $\hat{A}$-genus)
- $\pi_5(MSpin) = 0$
- $\pi_6(MSpin) = 0$
- $\pi_7(MSpin) = 0$
- $\pi_8(MSpin) = \mathbb{Z} \oplus \mathbb{Z}$ (From the Dirac operator and the Dirac operator coupled to the tangent bundle)

These invariants are largely determined by K-theory (via index theory) and mod 2 cohomology invariants (related to Stiefel-Whitney classes).

The Role of Anderson Duality: Torsion vs. Free Parts The two versions of the conjecture ($I_{\mathbb{Z}}$ and $I_{\mathbb{C}^*}$) are related by a fiber sequence of spectra:

$$I_{\mathbb{Z}} \to I_{\mathbb{C}} \to I_{\mathbb{C}^*}$$

This sequence induces a long exact sequence of homotopy groups for mapping spectra, which effectively separates the classification into free and torsion parts, analogous to the Universal Coefficient Theorem. For a spectrum E, this sequence is:

$$\cdots \to [E, \Sigma^m I_{\mathbb{Z}}] \to [E, \Sigma^m I_{\mathbb{C}}] \to [E, \Sigma^m I_{\mathbb{C}^*}) \to [E, \Sigma^{m-1} I_{\mathbb{Z}}] \to \dots$$

- $[E, \Sigma^m I_{\mathbb{C}}] \cong \mathrm{Hom}(\pi_m(E), \mathbb{C})$.
- The image of the map from $[E, \Sigma^m I_{\mathbb{Z}}]$ corresponds to the integer-valued invariants, i.e., homomorphisms $\pi_m(E) \to \mathbb{Z} \subset \mathbb{C}$.
- The kernel of the map from $[E, \Sigma^m I_{\mathbb{C}^*}]$ is the image from $[E, \Sigma^m I_{\mathbb{C}}]$. The cokernel, which maps to $[E, \Sigma^{m-1} I_{\mathbb{Z}}]$, is related to the torsion subgroup. Specifically, the kernel of this boundary map identifies the characters of the torsion subgroup of $\pi_{m-1}(E)$.

Thus, the $I_{\mathbb{Z}}$ formulation of the conjecture detects integer-valued invariants (like the integer quantum Hall effect), while the $I_{\mathbb{C}^*}$ formulation also captures torsion-valued invariants (phases protected by torsion elements in spin bordism groups).

2.2.3 Topological Field Theory Interpretation

The map from spin bordism to an Anderson dual spectrum can be understood as a fully extended, invertible **Topological Field Theory (TFT)**. This provides a powerful physical and geometric picture. The development of this perspective was deeply influenced by collaborations with M. Singer and conversations with P. Teichner and S. Stolz.

We can model the mapping space by considering a functor from a bordism category to a target category representing the algebraic data.

- For the $\mathbb{C}^*$ theory, the functor maps an n-manifold M^n to an object $Z(M^n) \in \mathbb{C}^*$, and an $(n+1)$-bordism N to a morphism $Z(N) \in \mathbb{C}$ such that if $M = \partial N$, then $Z(M) = e^{2\pi i Z(N)}$.
- For the $I_{\mathbb{Z}}$ theory, which is shifted by Σ^2, the target category is richer. It is modeled by the symmetric monoidal category of **super lines** (1-dimensional $\mathbb{Z}_2$-graded vector spaces). A functor from the 2-category of bordisms assigns a super line to an n-manifold, a map of super lines to an $(n+1)$-bordism, and an isomorphism of maps to an $(n+2)$-bordism with corners.

The physical picture is that a Hamiltonian system from $K(\mathbb{R}^n)$ defines a "material type." To evaluate the corresponding TFT on a manifold M^n, one "builds" the manifold out of this material (e.g., by placing the lattice system on it) and calculates its ground state. The lecturer provides a helpful analogy: "you want to make a dress out of that material... you cut your regions out of this material and then, you sew them together to make this picture." Since the phase is invertible, the ground state is expected to be one-dimensional, forming the super line L_M. The holonomy of this line bundle over a family of manifolds can recover the integer invariants.

2.2.4 Open Questions for Physics

This topological perspective raises fundamental questions about how these abstract structures are realized in concrete physical models.

1. **How do we recognize the Hamiltonian of a material on a generic manifold?** Given a physical system corresponding to a phase in $K(\mathbb{R}^n)$, what is the procedure for placing this system on an arbitrary curved spin manifold M^n? We are supposed to imagine the system is scaling under a renormalization flow to an IR limit. How can we look at a local patch on the manifold and confirm it is "made of" the correct material from $\mathbb{R}^n$? For specific models like the toric code or Kitaev wire, we have methods, but a general procedure for an arbitrary invertible phase is lacking.
2. **Can the space of all gapped phases be organized into a higher category?** The Kitaev conjecture addresses the *invertible* objects $K(\mathbb{R}^n)$. These should be the invertible objects within the larger collection of all

gapped phases $L(\mathbb{R}^n)$. This suggests a more general structure. Can $L(\mathbb{R}^n)$ be made into the objects of a **symmetric monoidal infinity n-category with duals**?

- **Objects:** Gapped phases on $\mathbb{R}^n$.
- **1-Morphisms:** Domain walls between phases (systems with a boundary).
- **2-Morphisms:** Defects or junctions between domain walls.
- etc.

If this structure exists, the theorem that the invertible objects form a spectrum would be one piece of a more general story, a physical realization of the Cobordism Hypothesis. A key challenge is to formulate a precise conjecture for what this higher category should be, extending Kitaev's conjecture beyond the invertible case, and to find guidance from the Hamiltonian lattice side about what a reasonable conjecture should be.

2.3 Lecture 3

2.3.1 Introduction: The Two Worlds and the Bridge

To paraphrase the military strategist Helmuth von Moltke, "No lecture plan survives first contact with the audience." This lecture series aims to build a bridge between two seemingly disparate fields: the mathematical world of **Homotopy Theory** and the physical world of **gapped Hamiltonian lattice systems**. The central pillar of this bridge is a conjecture by Alexei Kitaev, and by exploring this relationship, we hope to illuminate both sides and identify promising avenues for future research.

On the physics side, we consider the space of gapped, invertible, topological phases of matter. These systems, when organized by dimension, are conjectured to form the spaces of a spectrum in the sense of algebraic topology.

Definition 2.13. *The **Kitaev Spectrum**, denoted K, is a spectrum whose n-th space, K_n, is the space of phases of n-dimensional invertible, gapped Hamiltonian lattice models on $\mathbb{R}^n$. The path components, $\pi_0(K_n)$, classify these distinct phases.*

On the mathematical side, we have topological tools for classifying manifolds, chief among them being bordism theory. The Kitaev conjecture provides a stunning link between these two worlds.

Conjecture 2.4 (Kitaev)**.** *The Kitaev spectrum K is equivalent to a spectrum constructed from spin bordism theory, specifically the spectrum of maps into an Anderson dual of the integers:*

$$K \simeq \mathit{Maps}(MSpin, \Sigma^2 I_{\mathbb{Z}}) \quad \textit{or equivalently} \quad K \simeq \mathit{Maps}(MSpin, \Sigma I_{\mathbb{C}^*})$$

There is a natural map from the physical spectrum K to the simpler topological spectrum $\Sigma^2 I_{\mathbb{Z}}$, which is constructed via a universal property related to the fact that $\pi_0(K_0) \cong \mathbb{C}^*$. The deeper part of the conjecture is the existence of a lift to the full spin-equivariant spectrum, $\mathrm{Maps}(MSpin, \Sigma^2 I_{\mathbb{Z}})$. This lift is believed to be a consequence of fundamental physical principles.

Physical Principles $\implies$ *Mathematical Structure*

Reflection Positivity + Emergent Lorentz/Spin Symmetry $\implies$ Lift to Spin TQFT

2.3.2 A Strategy to Test the Conjecture

A powerful way to investigate this conjecture is to seek a counterexample. The strategy is as follows:

1. Identify a topological quantum field theory (TQFT) that is *framed* but provably *not spin*. In our diagram, this means finding an element κ in the space of framed TQFTs, $\pi_0(\mathrm{Maps}(S^0, \Sigma^2 I_{\mathbb{Z}}))$, that does not lift to the space of spin TQFTs, $\pi_0(\mathrm{Maps}(MSpin, \Sigma^2 I_{\mathbb{Z}}))$.

2. Attempt to construct a physical Hamiltonian lattice model for this specific TQFT κ.

3. Analyze the outcome:

   - If **YES**, a lattice model for κ can be constructed, then the Kitaev conjecture is false. We would have a physical system that corresponds to a non-spin TQFT.
   - If **NO**, then we must investigate the obstruction. Why can't we build a lattice model for this TQFT? The obstruction we encounter may reveal the physical mechanism (e.g., a necessary consequence of locality or unitarity) that enforces the emergent spin symmetry. This process will inform our understanding of the constraints that physics places on the space of possible TQFTs.

2.3.3 A Concrete Test Case: The 6D TQFT from $\pi_6(S^0)$

A prime candidate for such a test case 'kappa' arises from the classification of 5-dimensional lattice systems. According to the conjecture (ignoring the spin structure for a moment), the phases are classified by:

$$\pi_0(K_5) \cong \pi_0(\mathrm{Maps}(S^0, \Sigma^5\Sigma^2 I_{\mathbb{Z}})) \cong \mathrm{Hom}(\pi_6(S^0), \mathbb{Z})$$

The stable homotopy groups of spheres are fundamental invariants in topology. The first few are:

$$\pi_0 = \mathbb{Z}, \quad \pi_1 = \mathbb{Z}_2, \quad \pi_2 = \mathbb{Z}_2, \quad \pi_3 = \mathbb{Z}_{24}, \quad \pi_4 = 0, \quad \pi_5 = 0, \quad \pi_6 = \mathbb{Z}_2$$

The group we are interested in is $\pi_6(S^0) \cong \mathbb{Z}_2$. This group is generated by a class called ν^2. The corresponding TQFT, which we call κ, is a non-trivial invariant of 6-dimensional framed manifolds.

Crucially, this invariant does *not* lift to spin bordism. The relevant group in the spin case is $\pi_0(\mathrm{Maps}(MSpin, \Sigma^7 I_{\mathbb{Z}})) \cong \mathrm{Hom}(\pi_6(MSpin), \mathbb{Z})$, which is zero. This confirms that κ is a framed TQFT but not a spin TQFT.

This specific invariant is not merely a topological curiosity; it has appeared in theoretical physics. In Witten's work on the five-brane effective action in M-theory, this invariant arises as the Arf invariant of a quadratic function defined on the middle cohomology of a 6-manifold. The definition of this invariant requires the manifold to have more structure than just an orientation; it requires an integer lift of the fourth Stiefel-Whitney class, $W_4 = 0$, which is a stronger condition than a spin structure (which requires $W_1 = W_2 = 0$).

The challenge is then posed: can one construct a 5-dimensional Hamiltonian lattice model whose low-energy physics realizes this 6D Arf-invariant TQFT?

2.3.4 The Cobordism Hypothesis: A Framework for TQFTs

To properly discuss TQFTs and the structures involved, we need the language of higher categories, as provided by the Cobordism Hypothesis, a theorem formulated by Baez-Dolan and proven by Lurie.

Theorem 2.14 (The Cobordism Hypothesis)**.** *For any n, the* ***bordism ∞-n-category****, $Bord_n(point)$, is the* ***free symmetric monoidal ∞-n-category with duals*** *on a single object.*

This powerful statement organizes manifolds and their boundaries into a coherent categorical structure. For a general background space (or "tangential structure") B, the category Bord_n^B is defined as follows:

- **Objects** (0-morphisms): Closed, oriented 0-manifolds with a map to B.
- **1-Morphisms:** 1-dimensional bordisms between 0-manifolds.
- **k-Morphisms** ($k \leq n$): k-dimensional bordisms with corners.
- **Higher Morphisms** ($k > n$): Diffeomorphisms and paths between them. All morphisms for $k > n$ are invertible.

A key feature is that all objects are **fully dualizable**. This is a hierarchical duality property essential for defining TFTs on manifolds with boundaries.

The theorem implies that to define a symmetric monoidal functor $Z : \mathrm{Bord}_n \to \mathcal{C}$ (i.e., an n-dimensional TQFT with values in a target category $\mathcal{C}$), one only needs to specify the value of the functor on a single object: the point, $Z(\mathrm{pt}) \in \mathcal{C}$. The space of all such TQFTs is then equivalent to the space of fully dualizable objects in $\mathcal{C}$.

2.3.5 Connecting Physics and Topology: Mismatches and Open Questions

While the Cobordism Hypothesis provides a powerful mathematical framework, applying it to the physical world of lattice models reveals several subtle mismatches and deep questions.

Problem A: The Need for Non-Invertible Morphisms The Kitaev conjecture focuses on *invertible* TQFTs, which correspond to the invertible objects in the target category $\mathcal{C}$. However, the very structure of a symmetric monoidal category with duals relies on morphisms that are not isomorphisms. For example, the map that defines the dual of an object W is a morphism $W \otimes W^{\mathrm{dual}} \to \mathbf{1}$, which is certainly not invertible.

In physical terms, these non-invertible morphisms correspond to **domain walls**, **boundary conditions**, and higher-order **defects**. It seems that even to properly analyze the invertible phases, one must embed them in a larger structure that includes these non-invertible elements. This suggests that the definition

of the physical space $K(\mathbb{R}^n)$ should be extended to an ∞-n-category that incorporates domain walls as 1-morphisms, defects as 2-morphisms, and so on. Such a structure would naturally contain the spectrum of invertible phases as a consequence of the Cobordism Hypothesis.

Problem B: The Spin(n) vs. Spin(n+1) Mismatch A second, more technical issue concerns the precise spin structure required.

- The Kitaev conjecture requires a 'Spin(n)' structure to fully lift the map, corresponding to the symmetries of an n-dimensional system.
- However, theorems linking physics to TQFTs often rely on **reflection positivity**, a property of Euclidean field theory.

Theorem 2.15 (Freed-Hopkins)**.** *For an invertible TQFT, the property of being* ***reflection positive*** *implies that the theory is a* ***Spin(n+1)*** *fixed point.*

This gives a 'Spin(n+1)' structure, not the 'Spin(n)' structure needed for the conjecture. While for invertible theories a 'Spin(n+1)' action can be extended to a 'Spin' action, being a fixed point for the larger group does not automatically imply being a fixed point for the smaller one. This subtle gap was a key difficulty in extending these results beyond the invertible case.

There is significant ongoing work to bridge this gap. A large collaboration (including Ferrer, Krulewski, Muller, Penneys, Reutter, Scheimbauer, Stehouwer, and Johnson-Freyd, referred to as the "dagger gang") is developing the theory of higher dagger categories, which may provide the necessary tools to understand reflection positivity and duality in the non-invertible setting.

An Open Question for Physicists These mathematical considerations lead to a crucial question for the physics side. Even if we understand how reflection positivity provides a 'Spin(n+1)' structure on the objects (the phases), the morphisms (the domain walls) only carry an 'O(n)' action. This leaves many choices for what kinds of domain walls are physically permissible.

Open Question: Does the physical requirement of an emergent spin invariance in a lattice model place restrictions on the allowed types of domain walls between phases?

Answering this question is essential. The mathematicians need a "boundary condition" from the physicists to know which mathematical structures are the correct ones to model the full, extended nature of these physical theories.

2.4 Lecture 4

2.4.1 Introduction

This lecture marks the culmination of a journey exploring the deep and fruitful relationship between the physical world of **lattice models** and the mathematical framework of **topology**. I would like to begin by thanking the organizers for creating this unique opportunity for interdisciplinary dialogue. For those of us on the homotopy theory side, this conference has been invaluable. For a long time, many of us have been fascinated by these connections, sparked by conversations with physicists and mathematicians like Dan Freed, Constantine Teleman, Alexei Kitaev, Anton Kapustin, and Mike Freedman. However, the language and techniques of the lattice model community were often difficult to penetrate. A turning point for many was a talk by Daniel Spiegel, based on Peter Fieger's work, which made the algebraic quantum field theory approach accessible. This conference, bringing together pioneers like Peter and Bruno, has been a dream, finally making these foundational papers truly come alive.

Our central theme has been the relationship between physical systems and topological structures, guided by the Kitaev conjecture. This relationship is not a one-way street; each side provides crucial "boundary conditions on the imagination" for the other, inspiring new questions and revealing which abstract structures are physically relevant.

$$\textbf{Physics} \longleftrightarrow \textbf{Topology}$$

The goal of this final lecture is to explore the arrow going from right to left: how can we begin with an abstract topological field theory and construct a concrete lattice model? This "inverse problem" may shed light on some of the deepest physical principles, such as the origin of emergent Lorentz invariance.

2.4.2 The Grand Picture: A Recap

Let's briefly recap the overarching structure we've developed.

- **Lattice Models and Phases:** The starting point on the physics side is the universe of gapped Hamiltonian lattice models. These come in many flavors (bosonic, fermionic, with or without time-reversal symmetry, etc.).
- **Invertible Phases and Stable Homotopy:** For the special case of *invertible* phases, the Kitaev conjecture posits a direct link to stable homotopy theory. The spectrum of phases, K, is conjectured to be equivalent to a spectrum built from spin bordism: $K \simeq \mathrm{Maps}(MSpin, \Sigma^2 I_{\mathbb{Z}})$.
- **Non-Invertible Phases and Higher Categories:** The invertible phases should be viewed as the invertible objects within a much richer structure. The **Cobordism Hypothesis** organizes TQFTs into a hierarchy of *infinity n-categories*. It is natural to ask if the world of all gapped lattice models (including non-invertible ones) can be similarly organized, with domain walls as 1-morphisms, defects as 2-morphisms, and so on.

- **Physical Principles as Mathematical Lifts:** The conjecture that physical fermionic systems are described by *spin* TQFTs (i.e., that the map from K to TQFTs lifts from framed to spin bordism) is believed to be a consequence of two fundamental physical principles: **reflection positivity** and an **emergent Lorentz invariance**.

In the invertible case, the role of reflection positivity is reasonably well understood. For the more general non-invertible case, significant progress is being made by a large collaboration working on the theory of **dagger categories**. As reported by Theo Johnson-Freyd, this work appears to be quite well-developed and provides a potential topological model for reflection positivity. The origin of emergent Lorentz invariance, however, remains more mysterious.

2.4.3 The Inverse Problem: From TQFT to Lattice Model

Our main task today is to address the question: Given an abstract Topological Quantum Field Theory (TQFT), how can we construct a corresponding Hamiltonian lattice model?

A TQFT, which we'll denote by Z, is a symmetric monoidal functor from a bordism category to a target category. Crucially, it assigns algebraic data to manifolds. For an $(n+1)$-dimensional TQFT:

$$Z(\text{closed } n\text{-manifold}) = \text{a vector space}$$

This presents an immediate dimensional mismatch. A lattice model requires:

1. A local Hilbert space H_i at each lattice site i (a 0-dimensional object).
2. A local Hamiltonian term h_{ij} acting on neighboring Hilbert spaces, e.g., $H_i \otimes H_j$.

How can a theory that produces vector spaces from n-dimensional manifolds tell us what vector space to put at a 0-dimensional point?

The Role of Boundary Conditions The key idea, developed extensively by Freed and Teleman, is that one must choose a **boundary condition** for the TQFT. In the categorical language of the Cobordism Hypothesis, a TQFT is a functor $Z : \text{Bord}_{n+1} \to \mathcal{C}_{n+1}$, where $\mathcal{C}_{n+1}$ is the universal target category. A boundary condition is a choice of object in a lower-dimensional category that maps to the tensor unit of the higher category.

Let's consider a 2D theory for simplicity. The universal target category, $\mathcal{C}_2$, is the Morita category of superalgebras. To get a vector space, we can choose a boundary condition, which corresponds to picking a specific algebra object $A \in \mathcal{C}_2$. This choice effectively restricts the TQFT. The theory with this boundary condition becomes equivalent to the category of left A-modules.

This process gives us the fundamental data needed for a lattice model:

- **The Local Hilbert Space:** The vector space we assign to each lattice site is the underlying vector space of the algebra A itself.
- **The Hamiltonian:** The local Hamiltonian terms are constructed from the multiplication map of the algebra, $m : A \otimes A \to A$. This map defines a projection operator, and the Hamiltonian is typically taken to be $\mathbf{1} - P_{\text{ground}}$, resulting in a frustration-free model.

Thus, the "finite instruction set" for building the lattice model is the finite algebraic data of the chosen boundary condition (the algebra A).

Remark 2.16 (The Dagger Structure)**.** *A persistent issue is that TQFTs naturally produce complex vector spaces, whereas quantum mechanics requires Hilbert spaces. The promotion of a vector space to a Hilbert space requires a compatible Hermitian structure, often called a "dagger structure" in this context. While this is well-understood for invertible theories, the general non-invertible case relies on the developing theory of dagger categories. For this discussion, we will set aside this subtlety and assume such a structure can be chosen.*

The Tensor Network Perspective This construction process has a beautiful and computationally practical interpretation as the contraction of a **tensor network**.

1. **Step 1: Triangulate.** Take a manifold M on which you want to evaluate the theory, and choose a triangulation (or more generally, a handle decomposition).
2. **Step 2: Assign Tensors.** The algebraic data from the chosen boundary condition is used to define tensors. The algebra A itself can be represented as a tensor, and its multiplication map $m : A \otimes A \to A$ is another. These tensors are assigned to the vertices, edges, and faces of the triangulation.
3. **Step 3: Contract.** The value of the TQFT on the manifold M, $Z(M)$, is computed by contracting this entire tensor network—summing over all internal indices according to the connectivity of the triangulation.

Geometrically, the process of contracting the tensor network corresponds to performing surgery on the triangulation. Contracting along an edge corresponds to removing that edge and merging the adjacent simplices. By sequentially performing surgery on all internal parts of the triangulation, one simplifies the manifold down to a disjoint union of spheres, and the value of the TQFT is computed.

Examples and Obstructions This procedure has been successfully applied to well-known models:

- **AKLT Model:** For this model, one chooses the algebra $A = M_2(\mathbb{C})$, the 2×2 matrices. After accounting for an $SU(2)$ symmetry, this construction reproduces the AKLT Hamiltonian.

- **2D Unoriented TQFT:** A 2D unoriented theory can be constructed by choosing the algebra of quaternions, $A = \mathbb{H}$. This requires specifying additional data corresponding to how the unoriented nature is handled (isomorphisms $H \cong H^{\text{op}}$ and $H \cong H^{\text{dual}}$). With appropriate choices, the partition function on a closed surface M evaluates to $(-\lambda)^{\chi(M)}$, where $\chi(M)$ is the Euler characteristic. Reflection positivity constrains λ to be ± 1.

However, this construction is not universally applicable. The entire procedure hinges on the ability to choose a suitable boundary condition for the initial TQFT. The primary obstruction from the topological perspective is that an arbitrary TQFT is not guaranteed to admit such boundary conditions.

2.4.4 Conclusion and a Final Question

We have seen that there is a plausible, though highly intricate, path from an abstract TQFT to a concrete lattice model. The construction requires descending through the hierarchy of universal n-categories, choosing a boundary condition (an algebra), and interpreting its structure as the local data for a tensor network.

The upshot is that the ability to construct a lattice model appears to be deeply tied to the existence of a rich algebraic structure of boundary conditions, domain walls, and defects within the TQFT. This raises a final, crucial question that ties back to our original motivation. The Kitaev conjecture requires an emergent spin invariance in physical systems. We have seen that this is a subtle property, not guaranteed by the standard axioms of TQFT. This leads to the question:

Does the existence of a sufficiently rich structure of boundary conditions for a TQFT, when combined with a physical requirement like a Hermitian (dagger) structure, force the emergent spin invariance needed to satisfy the Kitaev conjecture?

This frames a precise mathematical question motivated by the physical problem. Answering it could provide the key to understanding the origins of Lorentz invariance and to finally, fully connecting the worlds of high-energy physics, condensed matter, and homotopy theory.

3 Pieter Naaijkens: Introduction to Superselection Sector Theory

Contents

3.1 Lecture 1

3.1.1 Introduction

Motivation and Goals The main issue we will consider is the classification of **gapped ground states** of quantum spin systems. Our approach will have the following features:

- We will always work in the **thermodynamic limit**, considering systems on an infinite lattice.
- We are interested in gapped ground states of **local Hamiltonians**.
- We want to classify these states according to some **equivalence relation**.
- Our focus will be on states exhibiting **topological order**, also known as long-range entanglement.

The central question is: can we find physically interesting **invariants** that allow us to distinguish and classify these states?

Why are these states interesting? States with topological order are particularly interesting because they can host **anyons**. These are quasi-particle excitations that generalize the familiar concepts of bosons and fermions. While exchanging two identical bosons leaves a state unchanged and exchanging two fermions introduces a minus sign, exchanging two anyons can result in more complex, non-trivial transformations. This phenomenon is known as **braided statistics**.

The algebraic properties of these anyons are described by a rich mathematical structure known as a **braided tensor C*-category** (in our examples, these will often be braided fusion or modular tensor categories). A key feature of these systems is their robustness: due to their topological nature, the properties of the anyons are stable against small, local perturbations of the system. This makes the category of anyons a powerful invariant for the phase of matter.

Our main goal is to answer the following question: Given a microscopic description of a system (e.g., a Hamiltonian), how can we systematically derive its category of anyons?

Our Approach The framework we will use is rooted in the **Doplicher-Haag-Roberts (DHR) theory** of superselection sectors, which originated in algebraic quantum field theory (AQFT). In AQFT, one associates algebras of observables to regions of spacetime. DHR theory was developed to classify the "charges" or superselection sectors in such theories.

- In (3+1) dimensions, DHR theory leads to Bose/Fermi statistics and culminates in the Doplicher-Roberts theorem, which relates symmetric tensor categories to representation categories of compact groups.
- In lower dimensions, the geometry allows for more interesting possibilities, namely the braided statistics of anyons.

Our approach is operator-algebraic, but on a technical level, it differs from classical AQFT as our local algebras are simple finite-dimensional matrix algebras.

3.1.2 The Toric Code

The toric code will serve as our guiding example. It is simple enough to be analyzed explicitly, yet it exhibits the non-trivial features we are interested in.

The Setup We consider an infinite square lattice $\mathbb{Z}^2$. The quantum degrees of freedom, or qubits, are located on the **edges** of this lattice.

- For each edge x, the local Hilbert space is $\mathcal{H}_x = \mathbb{C}^2$.
- The local observable algebra is $\mathcal{A}_x = M_2(\mathbb{C})$, the algebra of 2×2 complex matrices.

The observables are built from the **Pauli matrices**:

$$\sigma^x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma^y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma^z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

These matrices are self-adjoint, square to the identity, and anti-commute with each other: $\{\sigma^i, \sigma^j\} = 2\delta_{ij} I$. Together with the identity matrix, they form a basis for $M_2(\mathbb{C})$.

The Hamiltonian The dynamics are defined by a local, commuting-projector Hamiltonian. We define two types of local operators:

- For each vertex (or "star") s, the **star operator** is $A_s = \bigotimes_{j \in s} \sigma_j^x$.
- For each face (or "plaquette") p, the **plaquette operator** is $B_p = \bigotimes_{j \in p} \sigma_j^z$.

These operators all square to the identity ($A_s^2 = B_p^2 = I$) and they all commute with each other ($[A_s, A_{s'}] = [B_p, B_{p'}] = [A_s, B_p] = 0$). The local Hamiltonian for a finite region Λ is defined as:

$$H_\Lambda = \sum_{s \subset \Lambda} (I - A_s) + \sum_{p \subset \Lambda} (I - B_p).$$

Since all terms in the sum commute and $I - A_s, I - B_p$ are positive operators, the ground state will be one that is simultaneously an eigenstate of every A_s and B_p with eigenvalue 1.

The Ground State

Theorem 3.1. *The toric code on the infinite lattice has a unique* ***frustration-free*** *ground state, ω_0. This state is pure.*

Proof Sketch. A state is frustration-free if it minimizes the energy of every term in the Hamiltonian individually. Here, this means the state ω_0 must satisfy $\omega_0(A_s) = 1$ and $\omega_0(B_p) = 1$ for all stars s and plaquettes p. One can show that a state satisfying these conditions exists and is uniquely determined by them.

Using a lemma that a state uniquely determined by such conditions must be pure, we conclude ω_0 is pure. To show it is a ground state, we must verify the ground state condition, $-i\omega_0(A^*\delta(A)) \geq 0$. Since $\delta(A)$ is the commutator with $H = \sum(I - A_s) + (I - B_p)$, and ω_0 is invariant under multiplication by A_s and B_p, the condition simplifies to showing $\sum \omega_0(A^*(I - A_s)A) + \cdots \geq 0$, which is true since ω_0 is a state and the operators are positive. □

This unique state ω_0 will be our reference "vacuum" state. When the model is defined on a surface with non-trivial topology like a torus, the ground state becomes degenerate, which is the origin of the name "toric code" and its application in quantum error correction.

The GNS Representation of the Ground State The state ω_0 gives rise to a GNS representation $(\pi_0, \mathcal{H}_0, \Omega_0)$.

- The ground state vector Ω_0 is a "+1" eigenvector for all star and plaquette operators: $A_s\Omega_0 = \Omega_0$ and $B_p\Omega_0 = \Omega_0$. These are the **stabilizer conditions**.
- The GNS Hamiltonian H has Ω_0 as its zero-energy ground state ($H\Omega_0 = 0$) and is positive semi-definite ($H \geq 0$).
- The model is known to have a **spectral gap**: $\text{spec}(H) \cap (0, 2) = \emptyset$. This gap is stable against small local perturbations.

3.1.3 Excitations

Excitations above the ground state correspond to violations of the stabilizer conditions. These are our anyons.

Path Operators We can create excitations using string-like operators.

- A **path** ξ is a sequence of adjacent edges on the lattice. The associated operator is $F_\xi = \bigotimes_{j\in\xi} \sigma_j^z$.
- A **dual path** $\hat{\xi}$ is a path on the dual lattice. The operator is $F_{\hat{\xi}} = \bigotimes_{j\in\text{edges crossed by } \hat{\xi}} \sigma_j^x$.

These path operators commute with all A_s and B_p operators, *except* at the endpoints of the path.

- F_ξ anti-commutes with the star operators A_s at its endpoints.
- $F_{\hat{\xi}}$ anti-commutes with the plaquette operators B_p at its endpoints.

This means applying a path operator to the ground state creates a pair of excitations. For example, $H_\Lambda F_\xi \Omega_0 = 2N F_\xi \Omega_0$, where N is the number of endpoints of ξ inside Λ. The endpoints of F_ξ are called **electric charges**, and the endpoints of $F_{\hat{\xi}}$ are called **magnetic charges**.

Single Excitations and Localized Automorphisms To describe a single anyon, we work on the infinite lattice and send one endpoint of the path to

infinity. This is done by considering a semi-infinite path (or "ribbon") ξ that lies within some cone. The corresponding automorphism is defined as a limit:

$$\rho_\xi^Z(A) := \lim_{n\to\infty} \mathrm{Ad}_{F_{\xi_n}}(A) = \lim_{n\to\infty} F_{\xi_n} A F_{\xi_n}^*,$$

where ξ_n is the initial segment of the path ξ of length n. This limit is well-defined and defines an automorphism of the quasi-local algebra that is localized in the cone containing the path.

We can define three non-trivial types of such automorphisms, corresponding to the three types of anyons:

- ρ_ξ^Z (electric charge), created by a string of σ^z operators.
- $\rho_{\hat{\xi}}^X$ (magnetic charge), created by a string of σ^x operators.
- ρ^Y (dyon/fused charge), created by a combination of both.
- We also have the trivial automorphism $\rho^0 = \mathrm{id}$.

The state $\omega_0 \circ \rho_\xi^k$ describes the system with a single anyon of type k at the start point of the path ξ. While the automorphism ρ depends on the entire path, the resulting state $\omega_0 \circ \rho$ only depends on the endpoint. This is a key topological property.

Superselection Sectors We can now classify these single-anyon states.

Definition 3.2. *Two pure states ω_1, ω_2 are **equivalent** if their corresponding GNS representations are unitarily equivalent.*

A key result is that two pure states are inequivalent if and only if they can be distinguished by an observable "at infinity". That is, for any finite region Λ, there exists an observable A with support in the complement Λ^c such that $\omega_1(A) \neq \omega_2(A)$.

Theorem 3.3. *The four states corresponding to the different anyon types, $\omega_0 \circ \rho_x^k$ for $k \in \{0, X, Y, Z\}$, are mutually inequivalent.*

Proof Sketch. We can move the endpoint of the excitation locally, so we can assume they are all created at the same site x. To distinguish, say, an electric charge ($k = Z$) from a magnetic charge ($k = X$), we can use a "Wilson loop" operator. Consider a large closed dual loop $\hat{\xi}$ that encircles the site x. The operator $F_{\hat{\xi}}$ can be written as a product of all plaquette operators B_p inside the loop. The automorphism for the electric charge, ρ_x^Z, commutes with $F_{\hat{\xi}}$, so $\omega_0 \circ \rho_x^Z(F_{\hat{\xi}}) = \omega_0(F_{\hat{\xi}}) = 1$. However, the automorphism for the magnetic charge, ρ_x^X, anti-commutes with $F_{\hat{\xi}}$, so $\omega_0 \circ \rho_x^X(F_{\hat{\xi}}) = \omega_0(-F_{\hat{\xi}}) = -1$. Since we can make the loop arbitrarily large, we have found an operator at infinity that distinguishes the two states. A similar argument works for all pairs. □

These four inequivalent classes of states are the **superselection sectors** of the toric code. They correspond to the four types of anyons the model can host.

3.2 Lecture 2

3.2.1 The Story So Far

In the previous lecture, we considered the toric code on the infinite $\mathbb{Z}^2$ lattice. Our main achievements were:

- We constructed a unique, pure, frustration-free ground state ω_0.
- We defined automorphisms ρ^k for $k \in \{0, X, Y, Z\}$ that describe single anyons (trivial, electric, magnetic, and dyonic).
- We showed that these automorphisms give rise to four unitarily inequivalent classes of irreducible representations.

This raises two important questions:

1. How do we know that these are the "correct" or physically relevant representations to consider?
2. Does this set of representations possess additional structure, such as a fusion product corresponding to the fusion of anyons?

3.2.2 The Superselection Criterion

Superselection Rules The concept of a superselection rule dates back to the early days of quantum theory. It formalizes the idea that not all superpositions of states are physically realizable.

Definition 3.4 (Superselection Rule)**.** *Consider a representation $\pi : \mathfrak{A} \to \mathfrak{B}(\mathcal{H})$ and two unit vectors $\psi_1, \psi_2 \in \mathcal{H}$. We say that ψ_1 and ψ_2 satisfy a* ***superselection rule*** *if the state ω_θ corresponding to the superposition $\psi_\theta = (\psi_1 + e^{i\theta}\psi_2)/\sqrt{2}$ is independent of the relative phase θ. This can only happen if the representation π is not irreducible.*

Definition 3.5. *Two states ω_1 and ω_2 are called* ***not superposable*** *if in any representation π that contains vectors ψ_1, ψ_2 implementing these states, the state corresponding to a superposition is just an incoherent mixture:*

$$\omega_{\alpha\psi_1+\beta\psi_2}(A) = |\alpha|^2\omega_1(A) + |\beta|^2\omega_2(A)$$

for all $\alpha, \beta \in \mathbb{C}$ with $|\alpha|^2 + |\beta|^2 = 1$.

Theorem 3.6. *Two pure states ω_1, ω_2 are superposable if and only if their GNS representations are unitarily equivalent.*

This theorem provides the crucial link: inequivalent representations correspond to states that cannot be coherently superposed. We can think of these inequivalent classes as describing different "charges". An operator that is local (or even quasi-local) cannot change the total charge of a state. This is the physical reason why we cannot use local operators to remove a single anyon excitation from the ground state, and why such states are also ground states despite the local excitation energy.

The problem is that a C*-algebra generally has a vast number of inequivalent representations, most of which are unphysical. We need a criterion to select the interesting ones.

The DHR Superselection Criterion The Doplicher-Haag-Roberts (DHR) superselection criterion provides such a selection principle. It identifies representations that are "local perturbations" of a reference representation.

Definition 3.7 (Superselection Criterion)**.** *Let π_0 be an irreducible "reference" representation of $\mathfrak{A}$ (for us, the GNS representation of the ground state ω_0). A representation π satisfies the **superselection criterion** with respect to π_0 if for every cone $\Lambda \subset \mathbb{Z}^2$, the restriction of π to the algebra of observables in the complement of the cone, $\mathfrak{A}(\Lambda^c)$, is unitarily equivalent to the restriction of π_0:*

$$\pi \restriction_{\mathfrak{A}(\Lambda^c)} \cong \pi_0 \restriction_{\mathfrak{A}(\Lambda^c)} .$$

*An equivalence class of such representations is called a **superselection sector**.*

The interpretation is that sectors correspond to representations that are **localizable** (they look like the vacuum outside some cone) and **transportable** (we can move the cone of non-triviality around).

Theorem 3.8. *For the toric code, there are at least four irreducible superselection sectors.*

Proof Sketch. The four representations $\pi_k = \pi_0 \circ \rho^k$ that we constructed satisfy the criterion. For a given cone Λ containing the path defining ρ^k, the automorphism is the identity on $\mathfrak{A}(\Lambda^c)$, so π_k is equal to π_0 there. The transportability follows from the topological nature of the path operators. Since we already showed these four representations are inequivalent, they define four distinct sectors. It can be shown that these are, in fact, all the irreducible sectors. □

3.2.3 Towards a Tensor Category of Sectors

We now want to equip the set of superselection sectors with more structure, specifically a tensor product (fusion) and a braiding.

Monoidal Categories

Definition 3.9. *A **monoidal category** is a category $\mathcal{C}$ equipped with a bifunctor $\otimes : \mathcal{C} \times \mathcal{C} \to \mathcal{C}$ (the tensor product), a distinguished unit object $1_\mathcal{C}$, and families of natural isomorphisms:*

1. ***Associators:*** $\alpha_{a,b,c} : (a \otimes b) \otimes c \cong a \otimes (b \otimes c)$.
2. ***Unitors:*** $\lambda_a : 1_\mathcal{C} \otimes a \cong a$ *and* $\rho_a : a \otimes 1_\mathcal{C} \cong a$.

*These isomorphisms must satisfy consistency conditions known as the pentagon and triangle axioms. If the associators and unitors are all identity morphisms, the category is called **strict**.*

A Key Example: The Category of Endomorphisms Let $\mathfrak{A}$ be a unital C*-algebra. The category $\mathrm{End}(\mathfrak{A})$ has:

- **Objects:** Unital *-endomorphisms $\rho : \mathfrak{A} \to \mathfrak{A}$.
- **Morphisms:** For two endomorphisms ρ, σ, a morphism is an intertwiner $T \in \mathfrak{A}$ such that $T\rho(a) = \sigma(a)T$ for all $a \in \mathfrak{A}$.

This category becomes a strict monoidal category with the tensor product defined as composition:

- On objects: $\rho \otimes \sigma := \rho \circ \sigma$.
- On morphisms: For $S : \rho_1 \to \rho_2$ and $T : \sigma_1 \to \sigma_2$, their product is $S \otimes T := S\rho_1(T)$.

The Challenge for Superselection Sectors Our goal is to define a similar tensor product for our sectors. The anyon automorphisms ρ^k are objects in a category like $\mathrm{End}(\mathfrak{A})$. However, the morphisms are more complicated. The **charge transporters** V that relate representations with different paths, $V\pi_0 \circ \rho_x = \pi_0 \circ \rho'_x V$, are unitaries in $\mathfrak{B}(\mathcal{H}_0)$, but they are generally **not** in the image of the algebra, $\pi_0(\mathfrak{A})$. This is a problem, because the tensor product formula $S\rho_1(T)$ requires us to act with an endomorphism on an intertwiner, which is not well-defined if they live in different algebras.

Charge Transporters and von Neumann Algebras The solution lies in understanding the nature of the charge transporters. While the sequence of operators v_n used to construct the transporter does not converge in the norm topology of $\mathfrak{A}$, the sequence $\pi_0(v_n)$ *does* converge in the strong operator topology on $\mathfrak{B}(\mathcal{H}_0)$.

This leads us to consider von Neumann algebras.

Definition 3.10. *A unital *-subalgebra $\mathcal{M} \subset \mathfrak{B}(\mathcal{H})$ is a **von Neumann algebra** if it is equal to its double commutant, $\mathcal{M} = \mathcal{M}''$. Equivalently, $\mathcal{M}$ is closed in the weak (or strong) operator topology.*

Since $\pi_0(v_n)$ converges in the strong operator topology, its limit V must lie in the von Neumann algebra generated by the image of $\mathfrak{A}$, i.e., $V \in \pi_0(\mathfrak{A})''$.

Haag Duality and Localized Representations A crucial technical property is Haag duality.

Definition 3.11 (Haag Duality)**.** *A representation π_0 satisfies **Haag duality** for cones if for any cone Λ, the von Neumann algebra generated by observables in the cone equals the commutant of the algebra generated by observables in the complement: $\pi_0(\mathfrak{A}(\Lambda))'' = \pi_0(\mathfrak{A}(\Lambda^c))'$.*

This property holds for the toric code. It has important consequences:

1. **Localization of Intertwiners:** It implies that intertwiners between localized representations are themselves localized. If ρ_1, ρ_2 are localized in

cones $\Lambda_1, \Lambda_2 \subset \Lambda$, any intertwiner $v \in (\rho_1, \rho_2)$ must belong to the von Neumann algebra of the larger cone, $v \in \pi_0(\mathfrak{A}(\Lambda))''$.

2. **From Sectors to Localized Endomorphisms:** Any representation π satisfying the superselection criterion is unitarily equivalent to a representation ρ_Λ on the GNS space $\mathcal{H}_0$ which is localized in a cone Λ and maps observables in that cone into the corresponding von Neumann algebra: $\rho_\Lambda(\mathfrak{A}(\Lambda)) \subset \pi_0(\mathfrak{A}(\Lambda))''$.

This last point is key. While we don't get endomorphisms of $\mathfrak{A}$, we get something very close: endomorphisms of the von Neumann algebra $\pi_0(\mathfrak{A})''$. This will be sufficient to define the tensor product, which is the topic for the next lecture.

3.3 Lecture 3

3.3.1 The Monoidal Product

The Problem with the Tensor Product In the last lecture, we saw that the category $\mathrm{End}(\mathfrak{A})$ of endomorphisms of a C*-algebra $\mathfrak{A}$ has a natural strict monoidal product given by composition: $\rho \otimes \sigma = \rho \circ \sigma$. However, when we try to apply this to our category of superselection sectors, Δ_{DHR}, we run into a problem. The morphisms in our category (the intertwiners and charge transporters) are operators in the GNS Hilbert space, $v \in \mathfrak{B}(\mathcal{H}_0)$, and are generally not elements of the C*-algebra $\mathfrak{A}$ itself. Specifically, they lie in the von Neumann algebra generated by the image of $\mathfrak{A}$, e.g., $v \in \pi_0(\mathfrak{A}(\Lambda))''$.

This means that for an endomorphism ρ and an intertwiner v, the expression $\rho(v)$ is not well-defined, preventing us from defining the tensor product of morphisms $s \otimes t = s\rho_1(t)$.

The Auxiliary Algebra To solve this, we must extend our endomorphisms to a larger algebra that contains the intertwiners. This is done by constructing an **auxiliary algebra**.

1. Fix a "forbidden direction" by choosing a cone with angle $\theta \in [0, 2\pi)$ and opening $0 < \phi < \pi$.
2. Let $\mathcal{C}(\theta, \phi)$ be the set of all cones whose directions do not intersect the forbidden range $(\theta - \phi, \theta + \phi)$. This is a directed set whose union covers the plane.
3. We define the auxiliary algebra as the norm closure of the union of the von Neumann algebras associated with these cones:

$$\mathfrak{A}^{\mathrm{aux}} := \overline{\bigcup_{\Lambda \in \mathcal{C}(\theta,\phi)} \pi_0(\mathfrak{A}(\Lambda))''}^{\|\cdot\|}.$$

The original algebra $\mathfrak{A}$ is a subalgebra of $\mathfrak{A}^{\mathrm{aux}}$.

Lemma 3.12. *Let $\rho : \mathfrak{A} \to \mathfrak{B}(\mathcal{H}_0)$ be a localized and transportable representation (i.e., a sector). Then ρ can be extended to a *-homomorphism $\rho^a : \mathfrak{A}^{aux} \to \mathfrak{A}^{aux}$. This extension is weak-operator continuous on each cone algebra $\pi_0(\mathfrak{A}(\Lambda))''$.*

Proof Sketch. For any cone Λ in our directed set, we can choose a disjoint cone Λ'. By transportability, there is a unitary v and a representation ρ' localized in Λ' such that $\rho = \mathrm{Ad}_v \circ \rho'$. For any $a \in \mathfrak{A}(\Lambda)$, since Λ and Λ' are disjoint, $\rho'(a) = \pi_0(a)$. Thus, on this cone, $\rho(a) = v\pi_0(a)v^*$. Since conjugation by a unitary is weak-operator continuous, we can extend ρ from $\pi_0(\mathfrak{A}(\Lambda))$ to its weak closure, the von Neumann algebra $\pi_0(\mathfrak{A}(\Lambda))''$. By doing this consistently for all cones in the directed set, we obtain the full extension ρ^a. □

The DHR Category and its Tensor Product We can now give the final definition of our category of sectors.

Definition 3.13 (DHR Category)**.** *The category Δ_{DHR} has as objects the *-homomorphisms $\rho : \mathfrak{A} \to \mathfrak{A}^{aux}$ that are localized and transportable. The morphisms are the intertwiners in $\mathfrak{B}(\mathcal{H}_0)$.*

With the extended endomorphisms ρ^a, we can define the tensor product.

Definition 3.14 (Tensor Product)**.** *For objects $\rho, \sigma \in \Delta_{DHR}$ and morphisms $s \in (\rho_1, \rho_2)$, $t \in (\sigma_1, \sigma_2)$:*

- $\rho \otimes \sigma := \rho^a \circ \sigma$
- $s \otimes t := s\rho_1^a(t)$

Proposition 3.15. *With this definition, Δ_{DHR} is a strict monoidal C*-tensor category. The tensor unit is the identity map $\iota = id_{\mathfrak{A}}$, corresponding to the vacuum sector.*

The physical interpretation of this tensor product is the **fusion** of anyons: $\rho \otimes \sigma$ describes the sector obtained by creating an excitation of type σ and then an excitation of type ρ.

3.3.2 Braiding

A key achievement of the DHR framework is that it naturally yields a braiding, which describes the statistics of anyon exchange.

Definition 3.16. *A **braiding** on a monoidal category is a family of natural isomorphisms $\epsilon_{a,b} : a \otimes b \to b \otimes a$ that satisfy certain consistency conditions (the hexagon axioms). If $\epsilon_{b,a} \circ \epsilon_{a,b} = id$, the braiding is a **symmetry** (describing bosons/fermions).*

In general, for two endomorphisms $\rho, \sigma \in \mathrm{End}(\mathfrak{A})$, there is no reason for $\rho \circ \sigma$ and $\sigma \circ \rho$ to be related, so no braiding can be defined. However, the localization and transportability of our sectors allows us to construct one.

Construction of the Braiding

1. Take two sectors ρ, σ localized in cones $\Lambda_\rho, \Lambda_\sigma$.
2. Choose a new cone $\hat{\Lambda}_\sigma$ that is "to the left" of Λ_ρ.
3. By transportability, there is an equivalent representation $\hat{\sigma}$ localized in $\hat{\Lambda}_\sigma$, related by a charge transporter (a unitary intertwiner) $v \in (\sigma, \hat{\sigma})$.
4. Since Λ_ρ and $\hat{\Lambda}_\sigma$ are disjoint, the endomorphisms commute: $\rho \circ \hat{\sigma} = \hat{\sigma} \circ \rho$.
5. We can now define the braiding operator by moving σ to $\hat{\sigma}$, commuting, and moving back:
$$\epsilon_{\rho,\sigma} := v^*(\mathrm{id}_\rho \otimes v) = v^* \rho(v).$$
This defines an intertwiner from $\rho \otimes \sigma$ to $\sigma \otimes \rho$.

This construction is independent of the choices made (up to unitary equivalence) and defines a valid braiding on Δ_{DHR}.

Braiding in the Toric Code The explicit construction of the charge transporters in the toric code allows for an explicit calculation of the braiding. The key step is to compute $\rho^a(v)$. Since ρ^a is weak-operator continuous and v is a limit of local operators v_n, we have $\rho^a(v) = \lim_n \rho(v_n)$. For the toric code, the local path operators v_n are products of Pauli matrices, and one can compute that $\rho(v_n) = \pm v_n$. The sign depends on whether the string defining ρ crosses the string defining v_n. This leads to the braiding table:

ϵ_{ρ_1,ρ_2}	I	X	Y	Z
I	I	I	I	I
X	I	I	$-I$	$-I$
Y	I	$-I$	I	$-I$
Z	I	$-I$	$-I$	I

In particular, $\epsilon_{X,Z} \circ \epsilon_{Z,X} = (-I) \circ (-I) = -I$. This non-trivial sign shows that the electric (Z) and magnetic (X) charges are **abelian anyons**.

3.3.3 Fusion Rules and Direct Sums

The final piece of structure is the decomposition of tensor products into irreducible sectors.

Direct Sums and Subobjects To decompose objects, we need to be able to form direct sums and take subobjects. This requires the category to have additional linear structure. The category Δ_{DHR} is a C*-category, where Hom-sets are C*-algebras. This allows us to talk about projections, isometries, etc.

To define direct sums, we need to be able to find isometries v_1, v_2 in our algebra such that $v_1 v_1^* + v_2 v_2^* = I$. This is not always possible, but it is true if the underlying von Neumann algebra is "properly infinite".

Lemma 3.17. *For the toric code, the cone algebras $\pi_0(\mathfrak{A}(\Lambda))''$ are Type II_∞ factors, which are properly infinite.*

Corollary 3.18. *The category Δ_{DHR} for the toric code has direct sums and subobjects.*

This means we can decompose any object into a direct sum of subobjects corresponding to orthogonal projections in its endomorphism algebra. An object is **irreducible** (or simple) if its endomorphism algebra is trivial, $(\rho, \rho) = \mathbb{C}I$.

Fusion Rules The goal is to find the **fusion rules**, which describe the decomposition of the tensor product of two irreducible sectors ρ_i, ρ_j into a direct sum of other irreducible sectors ρ_k:

$$\rho_i \otimes \rho_j \cong \bigoplus_k N_{ij}^k \rho_k.$$

In general, this decomposition is not guaranteed to be finite. However, for the toric code, the situation is very simple. The fusion rules are those of the group $\mathbb{Z}_2 \times \mathbb{Z}_2$:

$$\rho^X \otimes \rho^X \cong \iota, \quad \rho^Z \otimes \rho^Z \cong \iota, \quad \rho^X \otimes \rho^Z \cong \rho^Y, \quad \text{etc.}$$

Physically, the rule $\rho^X \otimes \rho^X \cong \iota$ means that a state with two electric excitations has trivial total charge and thus belongs to the vacuum sector. The two anyons can annihilate each other.

3.4 Lecture 4

3.4.1 Further Structure: Duality and Rigidity

In the previous lectures, we have discussed superselection sectors, defined the monoidal product, decomposed these products, and defined the braiding. There is one more significant piece of structure to mention: **duality**, or **rigidity**. Because we are working with C*-categories, we can formulate this in a symmetric way.

Definition 3.19 (Conjugate)**.** *A **conjugate** for an object $\rho \in \Delta_{DHR}$ is a triple $(\overline{\rho}, r, \overline{r})$ with an object $\overline{\rho}$ and morphisms $r \in \mathrm{Hom}(id, \overline{\rho} \otimes \rho)$ and $\overline{r} \in \mathrm{Hom}(id, \rho \otimes \overline{\rho})$ such that the following **zig-zag equations** hold:*

$$(\overline{r}^* \otimes id_\rho)(id_\rho \otimes r) = id_\rho$$
$$(r^* \otimes id_{\overline{\rho}})(id_{\overline{\rho}} \otimes \overline{r}) = id_{\overline{\rho}}$$

These morphisms define a **conjugate charge**. Physically, the existence of $r \in \mathrm{Hom}(id, \overline{\rho} \otimes \rho)$ means that fusing the charge $\overline{\rho}$ with ρ yields a state containing the trivial sector (the vacuum). This structure allows for the definition of important concepts like the **quantum dimension** of a sector.

For the **toric code**, we can explicitly construct these conjugates; in fact, all sectors are self-conjugate since they square to the identity. In many other models, construction is also possible. However, it is not known if conjugates automatically exist in the general setting of quantum spin systems. In algebraic quantum field theory (AQFT), their existence can be guaranteed under certain conditions (like Poincaré covariance), but there are also examples of sectors in QFT for which a dual does not exist. In many algebraic descriptions of anyons, the existence of conjugates is often postulated. It would be valuable to find physical conditions that guarantee their existence in this context.

A brief note on a question that was raised: in the DHR approach for the toric code, one might construct representations involving infinite direct sums that do not have a dual. These are typically considered unphysical or "uninteresting" examples. The counter-examples to the existence of duals are generally of this "infinite type".

3.4.2 Completeness

A crucial question arises from our construction: have we found all the possible sectors? For the toric code, we constructed four irreducible sectors, as expected. But are there any other irreducible representations satisfying the superselection criterion that are not equivalent to one of these four? This is the question of **completeness**.

Approach 1: The Quantum Double

One powerful way to prove completeness is by relating our sectors to the representation theory of a known algebraic structure: the quantum double.

The **quantum double** $D(G)$ of a finite group G is a specific Hopf algebra built from the group algebra $\mathbb{C}[G]$. Its representation theory is well understood.

- The category of finite-dimensional representations, $\mathrm{Rep}_f D(G)$, is a **modular tensor category**.
- The irreducible representations (irreps) of $D(G)$ are in one-to-one correspondence with pairs (C, ρ), where C is a conjugacy class of G, and ρ is an irreducible representation of the centraliser of an element $g \in C$.

A recent result by Bols and Vadnerkar provides the crucial link to our superselection sectors.

Theorem 3.20 (Bols-Vadnerkar, arXiv:2310.19661)**.** *For the quantum double model for a finite group G, for each irreducible representation (C, ρ) of $D(G)$, there is a corresponding anyon sector $\pi^{(C,\rho)}$. The representations $\{\pi^{(C,\rho)}\}_{(C,\rho)}$ are pairwise disjoint (i.e., not unitarily equivalent), and any anyon sector is unitarily equivalent to one of them.*

In the context of my terminology, an "anyon sector" is an irreducible representation that satisfies the superselection criterion. For the toric code, the group is $G = \mathbb{Z}_2$. This group is abelian, so its conjugacy classes are just the individual elements, and the centralizers are the group itself. A quick calculation shows that there are exactly four such irreps. Since the theorem states that these are all the anyon sectors, we have indeed found all of them.

This leads to the main conclusion for the toric code model. Let Δ^f_{DHR} be the category of superselection sectors where we restrict to finite direct sums (or equivalently, where morphism spaces are finite-dimensional).

Theorem 3.21. *The category Δ^f_{DHR} for the toric code is braided tensor equivalent to $\mathrm{Rep}_f D(\mathbb{Z}_2)$.*

Proof sketch. We have already constructed representatives of the four sectors. The theorem by Bols and Vadnerkar establishes a one-to-one correspondence between these sectors and the irreps of $D(\mathbb{Z}_2)$. We can thus define a functor mapping our sectors to these irreps. We then check that this functor preserves the braiding and other tensor category structures. Since we have explicitly calculated all sectors and their properties, this becomes a matter of matching them up on both sides. □

This result confirms that our framework, built from first principles, gives the expected answer for this key example. The theory can also describe the twist and other structures that make the category modular.

Extension to Non-Abelian Models This analysis can be extended to quantum double models based on non-abelian finite groups G.

- If G is abelian, the analysis is very similar to the toric code.

- For non-abelian G, the "path operators" that create excitations are replaced by **ribbon operators**.
- A key difference is that irreps of $D(G)$ can have dimension $d > 1$. The corresponding ribbon operators then come in **multiplets** and are generally not unitary.
- This complicates the construction of the endomorphisms. Instead of conjugating by a single operator, it is easier to construct **amplimorphisms**, which are *-homomorphisms of the form $\rho : \mathfrak{A} \to M_d(\mathfrak{A})$. One can think of this as arranging the multiplet of ribbon operators into a $d \times d$ matrix and conjugating by that matrix.
- Using the property that the cone algebras are **properly infinite**, we can map these amplimorphisms back to endomorphisms of the auxiliary algebra $\mathfrak{A}^{aux}$, returning to our original framework. The cost is that this mapping is less explicit.
- Following this procedure, one can show that the resulting category of sectors is equivalent to $\mathrm{Rep}_f D(G)$, as expected. This confirms that abelian anyons correspond to automorphisms, while non-abelian anyons correspond to proper endomorphisms.

3.4.3 Stability of the Sector Category

So far, we have analyzed the sectors of a single, specific ground state. A fundamental question in the study of phases of matter is whether the structures we find are robust. Is the category of superselection sectors an **invariant of the quantum phase**? To answer this, we must first define what we mean by a phase. This is done by specifying an equivalence relation on states.

Definition 3.22 (Approximately Factorisable Automorphism (Informal))**.** *An automorphism $\alpha \in \mathrm{Aut}(\mathfrak{A})$ is **approximately factorisable** if for any cone Λ and $\delta > 0$, it can be decomposed as*

$$\alpha = \mathrm{Ad}(v) \circ \Xi \circ (\beta_\Lambda \otimes \beta_{\Lambda^c})$$

where:

1. *β_Λ and β_{Λ^c} are automorphisms localised inside Λ and its complement Λ^c, respectively.*
2. *Ξ is an automorphism acting only 'near the boundary' of Λ.*
3. *v is a unitary in the von Neumann algebra $\pi_0(\mathfrak{A})''$ that can be approximated by unitaries v_n localized in translates of slightly wider cones, with the error $||v_n - v||$ decaying sufficiently fast. A similar decomposition exists for α^{-1}.*

These automorphisms are generalizations of finite-depth quantum circuits. Crucially, they arise from physical situations. Consider a gapped local Hamiltonian H satisfying certain conditions (local topological order). If we add a sufficiently

small local perturbation Φ, the resulting Hamiltonian $H' = H+\Phi$ is also gapped. A key result in mathematical physics states that the ground states of H and H' are related by an approximately factorisable automorphism.

Definition 3.23 (Quantum Phase)**.** *Two pure states ω_1 and ω_2 are in the same **quantum phase** if there exists an approximately factorisable automorphism α such that $\omega_1 \circ \alpha = \omega_2$.*

Now we can state our central question more formally: If ω_1 and ω_2 are in the same phase, how are their corresponding DHR categories, $\Delta_{DHR}(\omega_1)$ and $\Delta_{DHR}(\omega_2)$, related?

Approximate Haag Duality

A technical challenge arises because an approximately factorisable automorphism α does not, in general, map a cone algebra to another cone algebra. This means that even if the GNS representation π_{ω_1} satisfies Haag duality, the perturbed representation $\pi_{\omega_2} \cong \pi_{\omega_1} \circ \alpha$ may not. We therefore need a weaker notion.

Definition 3.24 (Approximate Haag Duality)**.** *A representation π_0 satisfies **approximate Haag duality** if for every cone Λ and for small enough $\epsilon > 0$, there exists a unitary $U_{\Lambda,\epsilon}$ and a radius $R_\epsilon > 0$ such that*

$$\pi_0(\mathfrak{A}(\Lambda^c))' \subset U_{\Lambda,\epsilon}\pi_0(\mathfrak{A}((\Lambda_{-R_\epsilon})_\epsilon))''U^*_{\Lambda,\epsilon}$$

where $(\Lambda_{-R_\epsilon})_\epsilon$ denotes the cone Λ translated back by R_ϵ and with its opening angle widened by ϵ. Furthermore, the unitary $U_{\Lambda,\epsilon}$ can be approximated by unitaries in cone von Neumann algebras with sufficiently fast decay of the error.

Essentially, an operator commuting with everything outside a cone is no longer perfectly localized inside that same cone, but rather inside a slightly larger cone, up to a unitary conjugation. This unitary "sweeps up" the exponentially decaying tails created by the automorphism.

This weaker property has two crucial features:

1. If a state satisfies Haag duality, it also satisfies approximate Haag duality.
2. **Stability**: If a representation π satisfies approximate Haag duality, then so does $\pi \circ \alpha$ for any approximately factorisable automorphism α.

With this tool, the entire superselection theory can be reconstructed. The superselection criterion remains the same, but concepts like localization of representations and intertwiners now hold only up to decaying tails. The braiding becomes more complicated, requiring a limiting procedure to handle the tails. Despite these complexities, the framework remains intact. This leads to the main stability theorem.

Theorem 3.25 (Ogata, arXiv:2106.15741)**.** *Let ω_1 and ω_2 be pure gapped ground states of two uniformly bounded, finite-range interactions. Suppose they are in the same phase, i.e., $\omega_1 = \omega_2 \circ \alpha$ for some approximately factorisable*

automorphism α. If one (and hence both) of the GNS representations satisfies approximate Haag duality, then the corresponding DHR categories are unitarily braided monoidally equivalent.

Corollary 3.26. *The category of superselection sectors is an invariant of the quantum phase.*

3.4.4 Long-Range Entanglement and Trivial Sectors

The existence of anyons in the toric code is a manifestation of the state's **long-range entanglement**. We can make this connection precise within our framework.

Definition 3.27 (Product State)**.** *A pure state ω on $\mathfrak{A}$ is a **product state** with respect to a cone Λ if it is quasi-equivalent to a tensor product state, i.e., $\omega \approx \omega_\Lambda \otimes \omega_{\Lambda^c}$, where ω_Λ and ω_{Λ^c} are states on the algebras of the cone and its complement.*

This is a weaker notion than the usual site-by-site product state; we only require this structure with respect to a single partition of space into a cone and its complement.

Proposition 3.28. *Let ω be a pure product state with respect to a cone Λ. Then the von Neumann algebra $\pi_\omega(\mathfrak{A}(\Lambda))''$ is a Type I factor, and Haag duality holds for this cone: $\pi_\omega(\mathfrak{A}(\Lambda^c))' = \pi_\omega(\mathfrak{A}(\Lambda))''$.*

This has a profound consequence for the superselection structure.

Theorem 3.29 (Naaijkens, Ogata, arXiv:2102.07707)**.** *Let ω be a pure state which is a product state with respect to some cone Λ. Then the superselection theory with respect to π_ω is trivial, in the sense that every representation satisfying the superselection criterion is a (possibly infinite) direct sum of copies of π_ω.*

Proof Sketch. By the superselection criterion, any sector π can be localized to an endomorphism of the cone algebra $\pi_\omega(\mathfrak{A}(\Lambda))''$. But this algebra is a Type I factor (isomorphic to $\mathcal{B}(\mathcal{H})$ for some Hilbert space $\mathcal{H}$), which does not admit any "interesting" normal endomorphisms other than those corresponding to direct sums of the identity representation. □

This means that if a state is a product state, it only has the trivial anyon. We can now define long-range entanglement in this operator algebraic setting.

Definition 3.30 (Long-Range Entanglement)**.** *A state ω is **long-range entangled** if $\omega \circ \alpha$ is not a product state with respect to any cone for any approximately factorisable automorphism α.*

Combining the stability of the category and the triviality of sectors for product states, we arrive at a key physical insight.

Corollary 3.31. *If a state ω is not long-range entangled, the corresponding superselection structure $\Delta_{DHR}(\omega)$ is trivial.*

In other words, you need long-range entanglement to have anyons. This rigorously separates topologically ordered phases from trivial (or invertible) phases.

3.4.5 Completeness Revisited: A Subfactor Approach

The completeness proof based on the quantum double model relied on detailed knowledge of the specific model. An alternative approach exists, based on ideas from subfactor theory, which is more algebraic.

Consider two disjoint cones, Λ_A and Λ_B. We define two von Neumann algebras:

1. The algebra generated by observables in both cones: $\mathcal{R}_{AB} := \mathcal{R}_{\Lambda_A} \vee \mathcal{R}_{\Lambda_B}$.
2. The commutant of the algebra of the region outside both cones: $\hat{\mathcal{R}}_{AB} := \mathcal{R}'_{(\Lambda_A \cup \Lambda_B)^c}$.

From locality, it follows that $\mathcal{R}_{AB} \subset \hat{\mathcal{R}}_{AB}$. In general, this inclusion is strict, even if Haag duality holds. One can show that this forms an **irreducible subfactor**.

The key idea is that the larger algebra $\hat{\mathcal{R}}_{AB}$ is bigger precisely because it contains the **charge transporters**—the intertwiners between sectors localized in Λ_A and Λ_B, which are not necessarily contained in $\mathcal{R}_{AB}$.

We can quantify "how much bigger" $\hat{\mathcal{R}}_{AB}$ is using the **Jones index**, denoted $[\hat{\mathcal{R}}_{AB} : \mathcal{R}_{AB}]$. The index is ≥ 1, with equality if and only if the algebras are equal. If the index is finite, it roughly measures the relative size of the algebras.

Theorem 3.32 (Naaijkens, J. Math. Phys. 54 (2013))**.** *The number of irreducible superselection sectors is bounded from above by the index. If each sector has a conjugate, we have the stronger bound:*

$$\sum_i d(\rho_i)^2 \leq \mu_{AB} := \inf_{\Lambda_A, \Lambda_B} [\hat{\mathcal{R}}_{AB} : \mathcal{R}_{AB}]$$

where $d(\rho_i)$ is the quantum dimension of the sector ρ_i.

For abelian quantum double models, one can compute this index. For the toric code ($G = \mathbb{Z}_2$), the index is found to be $|G|^2 = 4$. Since we have explicitly constructed four sectors with $d(\rho_i) = 1$, we have $\sum d_i^2 = 1^2 + 1^2 + 1^2 + 1^2 = 4$. The bound is saturated, which implies that these must be all the sectors. This provides an alternative proof of completeness.

3.4.6 Summary

The main takeaways from this lecture series are:

- We can derive a **braided C*-category of superselection sectors** from first principles, starting from the local dynamics of a quantum spin system.

- This requires only a few general assumptions, namely a ground state gap and (approximate) Haag duality.
- Anyons are identified with representations satisfying the DHR superselection criterion: $\pi \restriction \mathfrak{A}(\Lambda^c) \cong \pi_0 \restriction \mathfrak{A}(\Lambda^c)$ for any cone complement Λ^c.
- The category can be constructed explicitly in models like the toric code or quantum double models, yielding the expected results (e.g., $\mathrm{Rep}_f D(G)$).
- The category is a robust **invariant of the quantum phase**, stable under small, local perturbations of the dynamics.
- A non-trivial superselection structure requires the ground state to have **long-range entanglement**.

4 Alexei Kitaev: Local Definitions of Gapped Hamiltonians and Invertible States

Contents

4.1 Lecture 1

4.1.1 Introduction and Conjectures

The lecture begins with conjectures regarding the classification of quantum phases of matter. The central idea is the existence of Ω-spectra, denoted B for bosonic systems and F for fermionic systems, which describe states of spin and fermionic Hamiltonians, respectively.

For some integer n, one can consider a construction B_n on a sphere S^n relative to a point. For $n = 0$, this space is conjectured to be:

$$B_0 = \mathbb{C}P^\infty$$

Similarly, for fermionic systems, there is a corresponding spectrum F_n, with:

$$F_0 = \mathbb{C}P^\infty \times \mathbb{Z}_2$$

With the help of Michael Hopkins and Daniel Freed, it was further conjectured that these spectra might be related to mapping spaces from cobordism spectra to the integer Eilenberg-MacLane spectrum $I\mathbb{Z}$:

$$B \stackrel{?}{=} \text{Maps}(MSO, I\mathbb{Z}(2))$$
$$F \stackrel{?}{=} \text{Maps}(MSpin, I\mathbb{Z}(2))$$

The shift by 2 in $I\mathbb{Z}(2)$ is necessary because the second homotopy group of the target is $\mathbb{Z}$. While progress on these specific conjectures has been limited without strong assumptions like Lorentz symmetry and reflection positivity, they motivate a general program for studying gapped phases.

The general program proceeds in steps:

1. Start with **concrete objects**, such as ground states of gapped Hamiltonians.
2. Impose conditions like **invertibility** to abstract these into a mathematical structure.
3. This should lead to an **abstract Ω-spectrum**.
4. Finally, one would calculate this to get a **concrete spectrum**.

This lecture focuses on the first and last steps, illustrated by the example of free fermions, where the final concrete spectrum is that of KO-theory. The key principles guiding the necessary definitions are **locality** (properties are determined by small regions) and **softness** (definitions are robust under small, arbitrary deformations).

4.1.2 The Transverse-Field Ising Model (TFIM)

We begin our concrete analysis with a canonical model in condensed matter physics: the one-dimensional transverse-field Ising model (TFIM).

Setup and Hamiltonian Consider a chain of N spin-$\frac{1}{2}$ particles (qubits). The Hilbert space for each spin is $\mathbb{C}^2$, and the total Hilbert space is $(\mathbb{C}^2)^{\otimes N}$. We use the Pauli matrices to describe operators on each site ℓ:

$$\sigma^x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma^y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma^z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

An operator acting on site ℓ is written as $\sigma_\ell^\alpha = I \otimes \cdots \otimes \sigma^\alpha \otimes \cdots \otimes I$, where σ^α is at the ℓ-th position.

Definition 4.1 (TFIM Hamiltonian)**.** *The Hamiltonian for the transverse-field Ising model on a chain of N sites is given by:*

$$H = -h \sum_{i=1}^{N} \sigma_i^z - J \sum_{i=1}^{N-1} \sigma_i^x \sigma_{i+1}^x$$

Here, h is the strength of the transverse magnetic field, and J is the strength of the ferromagnetic nearest-neighbor interaction. We consider an open chain without periodic boundary conditions for now.

We define the spin-up and spin-down basis states as:

$$|\uparrow\rangle \equiv |0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad |\downarrow\rangle \equiv |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

These are the eigenstates of σ^z. We also use the eigenstates of σ^x:

$$|\rightarrow\rangle = \frac{|\uparrow\rangle + |\downarrow\rangle}{\sqrt{2}}, \quad |\leftarrow\rangle = \frac{|\uparrow\rangle - |\downarrow\rangle}{\sqrt{2}}$$

Limiting Cases and Perturbation Theory We can understand the physics of the model by examining its limiting cases.

Case 1: No Interactions ($J = 0, h > 0$) When $J = 0$, the Hamiltonian is simply $H = -h \sum_{i=1}^{N} \sigma_i^z$. The spins are decoupled. The ground state is the one where all spins align with the field, i.e., all spins are up:

$$|\psi_0\rangle = |\uparrow\uparrow \ldots \uparrow\rangle$$

The ground state energy is $E_0 = -Nh$. An elementary excitation consists of flipping a single spin at site j.

$$|\psi_j\rangle = |\uparrow \ldots \uparrow \underbrace{\downarrow}_{j} \uparrow \ldots \uparrow\rangle$$

The energy of this state is $E_j = -Nh + 2h$. The energy cost to create one such excitation (a "quasi-particle") is $2h$. These excitations are static; a flipped spin at site j remains at site j.

Case 2: Ferromagnetic Limit ($h = 0, J > 0$) When $h = 0$, the Hamiltonian is $H = -J\sum_{i=1}^{N-1} \sigma_i^x \sigma_{i+1}^x$. To minimize the energy, all adjacent spins must have the same orientation in the x-direction. This leads to two degenerate ground states:

$$|\rightarrow\rightarrow\cdots\rightarrow\rangle \quad \text{and} \quad |\leftarrow\leftarrow\cdots\leftarrow\rangle$$

The elementary excitations in this case are **domain walls**, which are boundaries between regions of oppositely oriented spins, e.g., $|\rightarrow\rightarrow|\leftarrow\leftarrow\rangle$.

Case 3: Perturbative Regime ($J \ll h$) Let's consider a small interaction term J as a perturbation to the $J = 0$ case. The full Hamiltonian is $H = H^{(0)} + V$, where $H^{(0)} = -h\sum \sigma_i^z$ and $V = -J\sum \sigma_i^x \sigma_{i+1}^x$.

The single-spin-flip states $\{|\psi_j\rangle\}$ are degenerate eigenstates of $H^{(0)}$ with energy $E_p = E_0 + 2h$. The perturbation V mixes these states. The operator $\sigma_i^x \sigma_{i+1}^x$ flips spins at sites i and $i+1$. When it acts on a state with one flipped spin, say at site j, it can move the flipped spin. For example, the term $\sigma_{j-1}^x \sigma_j^x$ acting on $|\psi_j\rangle$ produces a state proportional to $|\psi_{j-1}\rangle$.

The effective Hamiltonian acting on the subspace of single quasi-particle states is called the quasi-particle Hamiltonian, H_{sp}. We find its matrix elements by projecting H onto this subspace.

$$\langle\psi_k|H|\psi_j\rangle = (E_0 + 2h)\delta_{kj} + \langle\psi_k|V|\psi_j\rangle$$

After subtracting the ground state energy E_0, the effective Hamiltonian for the quasi-particle at site j is:

$$H_{sp}|\psi_j\rangle = (2h)|\psi_j\rangle - J|\psi_{j-1}\rangle - J|\psi_{j+1}\rangle$$

This describes a particle hopping on a 1D lattice. For a periodic chain of length N, the eigenvectors are plane waves:

$$|\psi(q)\rangle = \sum_{j=1}^{N} e^{iqj}|\psi_j\rangle, \quad \text{with } q = \frac{2\pi k}{N} \text{ for } k \in \{0, \ldots, N-1\}$$

The corresponding eigenvalues give the energy dispersion of the quasi-particles:

$$E(q) = 2h - 2J\cos(q)$$

For a chain with open boundary conditions, the eigenvectors are standing waves, e.g., $\propto \sin(qL)$, with quantized momenta $q = \frac{\pi k}{N+1}$.

The quasi-particles are now dynamic; they propagate with a group velocity $v_g = \frac{\partial E(q)}{\partial q} = 2J\sin(q)$.

4.1.3 Exact Solution via Jordan-Wigner Transformation

The TFIM is special because it can be solved exactly for any h and J. The key is the Jordan-Wigner (JW) transformation, which maps the spin operators to fermionic operators.

Definition 4.2 (Majorana Fermion Operators)**.** *We introduce a pair of operators, $C_{2\ell-1}$ and $C_{2\ell}$, for each spin site $\ell = 1, \ldots, N$. They are defined in terms of the spin operators as follows:*

$$C_{2\ell-1} = \left(\prod_{k=1}^{\ell-1} \sigma_k^z\right) \sigma_\ell^x$$
$$C_{2\ell} = \left(\prod_{k=1}^{\ell-1} \sigma_k^z\right) \sigma_\ell^y$$

The string of σ_k^z operators is non-local. These operators C_j are Hermitian ($C_j^\dagger = C_j$), square to the identity ($C_j^2 = 1$), and anti-commute for different indices ($C_j C_k = -C_k C_j$ for $j \neq k$).

Proposition 4.3 (Clifford Algebra)**.** *The Majorana operators $\{C_j\}_{j=1}^{2N}$ are generators of a real Clifford algebra, satisfying the relation:*

$$C_j C_k + C_k C_j = 2\delta_{jk} I$$

We can visualize this mapping by replacing each spin site ℓ with a box containing two "Majorana modes," $C_{2\ell-1}$ and $C_{2\ell}$.

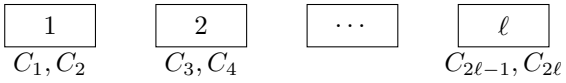

In terms of these Majorana operators, the TFIM Hamiltonian becomes quadratic. We need the following identities:

- $\sigma_\ell^z = -i\sigma_\ell^x \sigma_\ell^y = -i\left(\prod_{k=1}^{\ell-1} \sigma_k^z\right)^{-1} C_{2\ell-1} \left(\prod_{k=1}^{\ell-1} \sigma_k^z\right)^{-1} C_{2\ell} = -iC_{2\ell-1}C_{2\ell}$.
- $\sigma_\ell^x \sigma_{\ell+1}^x = C_{2\ell-1}(\prod_k \sigma_k^z)^{-1}(\prod_k \sigma_k^z)C_{2\ell+1} = -iC_{2\ell}C_{2\ell+1}$. (A careful derivation gives $\sigma_\ell^x \sigma_{\ell+1}^x = \sigma_\ell^x(\sigma_\ell^z)\sigma_{\ell+1}^y = (-i\sigma_\ell^y)\sigma_{\ell+1}^y = \cdots = -iC_{2\ell}C_{2\ell+1}$)

The Hamiltonian becomes:

Theorem 4.4 (Majorana Chain Hamiltonian)**.** *The TFIM Hamiltonian is equivalent to a quadratic Hamiltonian of Majorana fermions, often called the Majorana chain:*

$$H = ih \sum_{i=1}^{N} C_{2i-1}C_{2i} + iJ \sum_{i=1}^{N-1} C_{2i}C_{2i+1}$$

This describes nearest-neighbor "hopping" terms for Majorana fermions. The h term couples Majoranas within a site, and the J term couples Majoranas between adjacent sites.

4.1.4 Quadratic Fermion Hamiltonians

The Majorana form of the TFIM is an instance of a general class of solvable models.

Definition 4.5 (General Quadratic Hamiltonian)**.** *A general quadratic Hamiltonian for Majorana fermions can be written as:*

$$H(A) = \frac{i}{4} \sum_{j,k=1}^{2N} A_{jk} C_j C_k$$

where A is a $2N \times 2N$ real, anti-symmetric matrix.

For the TFIM, the matrix A has a block tri-diagonal structure. For a 2×2 block corresponding to sites i and $i+1$, the matrix of coefficients is proportional to:

$$A \propto \begin{pmatrix} 0 & 2h & 0 & 0 \\ -2h & 0 & 2J & 0 \\ 0 & -2J & 0 & 2h \\ 0 & 0 & -2h & 0 \end{pmatrix}$$

The operators $-iH(A)$ form a representation of the Lie algebra $\mathfrak{so}(2N)$:

$$[-iH(A), -iH(B)] = -iH([A, B])$$

where $[A, B]$ is the matrix commutator.

Diagonalization and Ground State Any real anti-symmetric matrix A can be brought to a canonical block-diagonal form by an orthogonal transformation $Q \in O(2N)$:

$$Q^T A Q = \bigoplus_{\nu=1}^{N} \begin{pmatrix} 0 & E_\nu \\ -E_\nu & 0 \end{pmatrix}, \quad E_\nu \geq 0$$

The values $\{E_\nu\}$ are related to the energies of the fermionic modes.

To find the ground state, we define a set of annihilation operators. A standard choice is:

$$a_\ell = \frac{C_{2\ell-1} + iC_{2\ell}}{2}, \quad a_\ell^\dagger = \frac{C_{2\ell-1} - iC_{2\ell}}{2}$$

The "bare" vacuum $|0\rangle$ is the state annihilated by all a_ℓ: $a_\ell|0\rangle = 0$ for all ℓ. This vacuum corresponds to a choice of **complex structure**, an operator B with $B^2 = -1$ that partitions the Majorana operators into annihilation and creation parts. For the standard choice above, B is the matrix mapping $C_{2\ell-1} \to C_{2\ell}$ and $C_{2\ell} \to -C_{2\ell-1}$.

The Hamiltonian $H(A)$ is generally not diagonal in this basis. The orthogonal matrix Q defines a new basis of Majorana operators $\tilde{C}_j = \sum_k Q_{kj} C_k$. In this new basis, the Hamiltonian takes the block-diagonal form:

$$H = \frac{i}{4} \sum_{\nu=1}^{N} E_\nu (\tilde{C}_{2\nu-1} \tilde{C}_{2\nu} - \tilde{C}_{2\nu} \tilde{C}_{2\nu-1}) = \frac{i}{2} \sum_{\nu=1}^{N} E_\nu \tilde{C}_{2\nu-1} \tilde{C}_{2\nu}$$

We can then define a new set of annihilation operators ("quasi-particle operators"):

$$\tilde{a}_\nu = \frac{\tilde{C}_{2\nu-1} + i\tilde{C}_{2\nu}}{2}$$

The ground state of the Hamiltonian $H(A)$, denoted $|\psi_A\rangle$, is the state annihilated by all the new operators:

$$\tilde{a}_\nu|\psi_A\rangle = 0 \quad \text{for all } \nu = 1, \ldots, N$$

This state $|\psi_A\rangle$ is the true ground state of the TFIM. The operators $\tilde{a}_\nu^\dagger$ create quasi-particle excitations above this ground state with energy E_ν. The full spectrum of the TFIM is thereby determined.

4.2 Lecture 2

4.2.1 Diagonalization of Quadratic Majorana Hamiltonians: A Recap

We begin by recalling the general procedure for solving a quadratic Majorana Hamiltonian of the form $H(A) = \frac{i}{4}\sum_{j,k} A_{jk}C_jC_k$, where A is a real, anti-symmetric $2N \times 2N$ matrix.

The algorithm is as follows:

1. **Diagonalize A:** Since A is real and anti-symmetric, its eigenvalues are purely imaginary and come in conjugate pairs $\pm i\epsilon_\ell$ (with $\epsilon_\ell \geq 0$). We can find an orthogonal matrix $Q \in O(2N)$ such that A is brought to a block-diagonal form:

$$Q^T A Q = \bigoplus_{\ell=1}^{N} \begin{pmatrix} 0 & \epsilon_\ell \\ -\epsilon_\ell & 0 \end{pmatrix}$$

   The columns of Q are constructed from the eigenvectors of A. If $AU_\ell = i\epsilon_\ell U_\ell$, then U_ℓ is a complex vector, and its real and imaginary parts can be used to form two orthogonal columns of Q.

2. **Define New Operators:** We use Q to define a new basis of Majorana operators:

$$\tilde{C}_j = \sum_{k=1}^{2N} (Q^T)_{jk} C_k = \sum_{k=1}^{2N} Q_{kj} C_k$$

   From these, we form a new set of fermionic annihilation operators (quasi-particle operators):

$$\tilde{a}_\ell = \frac{\tilde{C}_{2\ell-1} + i\tilde{C}_{2\ell}}{2}$$

3. **Solve the System:** The Hamiltonian expressed in terms of these new operators is diagonal:

$$H(A) = \sum_{\ell=1}^{N} \frac{\epsilon_\ell}{2}(2\tilde{a}_\ell^\dagger \tilde{a}_\ell - 1) = \sum_{\ell=1}^{N} \epsilon_\ell \left(\tilde{n}_\ell - \frac{1}{2}\right)$$

   where $\tilde{n}_\ell = \tilde{a}_\ell^\dagger \tilde{a}_\ell$ is the number operator for the ℓ-th quasi-particle mode. The ground state (new vacuum) $|\tilde{\Psi}_0\rangle$ is the state annihilated by all $\tilde{a}_\ell$. The excited states are created by applying $\tilde{a}_\ell^\dagger$ to the ground state. The energy of an excitation is ϵ_ℓ.

4.2.2 Application to the Majorana Chain

Let's apply this to the Majorana chain Hamiltonian, which is $H(A)$ for a specific A representing nearest-neighbor couplings.

$$H = ih\sum_{i=1}^{N} C_{2i-1}C_{2i} + iJ\sum_{i=1}^{N-1} C_{2i}C_{2i+1}$$

If we impose periodic boundary conditions (coupling site N to site 1), the system becomes translationally invariant. We can use a Fourier transform to diagonalize the matrix A. The eigenvectors take the form of plane waves, characterized by a momentum q. For each q, the problem reduces to diagonalizing a 2×2 matrix $A(q)$. The spectrum of this matrix gives the quasi-particle energy dispersion $\epsilon(q)$.

The resulting dispersion relation is:

$$\epsilon(q) = 2|h - Je^{iq}| = 2\sqrt{(h - J\cos q)^2 + (J\sin q)^2}$$

The system has an energy gap (all $\epsilon(q) > 0$) unless the expression under the square root vanishes. This happens only when $q = 0$ and $|h| = |J|$. The line $|h| = |J|$ is a phase boundary separating two distinct phases.

Boundary Effects and Majorana Zero Modes The analysis above for the periodic chain reveals a subtlety. The Transverse-Field Ising Model (TFIM) in the limit $h = 0, J > 0$ has two degenerate ground states ($|\rightarrow\cdots\rightarrow\rangle$ and $|\leftarrow\cdots\leftarrow\rangle$). However, our Majorana chain solution with periodic boundary conditions has a unique ground state unless $|h| = |J|$.

The discrepancy arises from the non-local nature of the Jordan-Wigner transformation. A periodic spin chain does not map to a simple periodic Majorana chain. If we instead consider the original open-boundary chain, the picture becomes consistent.

Let's analyze the two limiting cases for the open chain:

1. **Trivial Phase ($J = 0, h > 0$):** The Hamiltonian only contains terms $ihC_{2i-1}C_{2i}$. The Majorana operators are paired up within each site. All modes are gapped with energy $2h$. The ground state is unique.

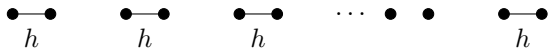

2. **Topological Phase ($h = 0, J > 0$):** The Hamiltonian contains terms $iJC_{2i}C_{2i+1}$. The Majorana operators are paired between adjacent sites.

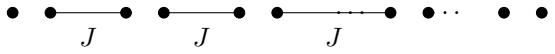

   In this configuration, the two Majorana operators at the ends of the chain, C_1 and C_{2N}, are left unpaired. They do not appear in the Hamiltonian.

These two operators combine to form a single fermionic mode $\tilde{a}_0 = (C_1 + iC_{2N})/2$ which has exactly zero energy ($\epsilon_0 = 0$). This is a **Majorana zero mode**.

The existence of this zero mode means the ground state is two-fold degenerate: the mode can be empty ($n_0 = 0$) or occupied ($n_0 = 1$) at no energy cost. This correctly reproduces the two-fold degeneracy of the ferromagnetic Ising model. The presence of such protected, localized boundary modes is the hallmark of a topological phase.

4.2.3 The Big Picture: A Hierarchy of Gapped Phases

The Majorana chain provides a concrete example that fits into a larger classification scheme for phases of matter. We can think of a hierarchy of classes of systems, each contained within the next:

$$\text{Free Fermion} \subseteq \text{Invertible} \subseteq \text{Topological} \subseteq \text{Gapped (in the bulk)}$$

The defining features that distinguish these classes are:

- **Gapped (in the bulk):** The most general class. Excitations in the bulk of the system cost a finite amount of energy, but the system may have complex non-universal properties. Examples include systems with **fractons**.
- **Topological:** These phases are robust against local perturbations and deformations of the geometry. They are characterized by universal properties. They may host **anyons** (quasi-particles with non-trivial braiding statistics, like in the Toric Code) and/or protected **ungappable boundary modes**.
- **Invertible:** A special class of topological phases that do not have anyons. They are "invertible" in the sense that for any system, there exists an "inverse" system such that their combination is trivial (can be continuously deformed to a simple product state without closing the bulk energy gap). They may still have ungappable boundary modes (e.g., the Integer Quantum Hall Effect or our topological Majorana chain).
- **Free Fermion:** Systems described by quadratic Hamiltonians of fermions. These are necessarily invertible.

4.2.4 Invertibility in Zero Dimensions

Let's make the definition of invertibility precise, starting with the simple case of 0-dimensional systems (i.e., just a single matrix A, not an extended system).

Definition 4.6 (Invertibility for $d = 0$)**.** *A gapped free fermion system described by a real, anti-symmetric matrix A is **invertible** if the combined system $A \oplus (-A)$ can be continuously connected to a trivial gapped system without closing the energy gap.*

Proof. Let the initial system be $A \oplus (-A)$, which acts on the space $M \oplus M$. This can be written as the matrix $A_0 = \begin{pmatrix} A & 0 \\ 0 & -A \end{pmatrix}$. Let the trivial system be $A_1 = \begin{pmatrix} 0 & A \\ -A & 0 \end{pmatrix}$. We want to find a path $A(\theta)$ for $\theta \in [0, \pi/2]$ that connects A_0 to A_1 while remaining gapped. Consider the path:

$$A(\theta) = \cos(\theta)A_0 + \sin(\theta)A_1 = \begin{pmatrix} \cos(\theta)A & \sin(\theta)A \\ -\sin(\theta)A & -\cos(\theta)A \end{pmatrix}$$

To check the gap, we compute $-A(\theta)^2$:

$$-A(\theta)^2 = -\begin{pmatrix} (\cos^2\theta - \sin^2\theta)A^2 & 0 \\ 0 & (\cos^2\theta - \sin^2\theta)A^2 \end{pmatrix} = \cos(2\theta)\begin{pmatrix} -A^2 & 0 \\ 0 & -A^2 \end{pmatrix}$$

For $\theta \in [0, \pi/4]$, $\cos(2\theta) > 0$. If the original system is gapped, i.e., $-A^2 \geq \Delta^2 I > 0$, then $-A(\theta)^2 \geq \cos(2\theta)\Delta^2 I$, so the gap remains open. A similar path connects A_1 to a completely trivial matrix. Thus, $A \oplus (-A)$ is deformable to a trivial state. □

4.2.5 A Framework for Locality

To extend these ideas to realistic, spatially extended systems, we need a rigorous way to handle locality.

Local Matrices Let $S \subseteq \mathbb{R}^n$ be a set of sites. For each site $x \in S$, we associate a finite-dimensional Euclidean space M_x. The total space is $M = \bigoplus_{x \in S} M_x$.

Definition 4.7 (Local Matrix)**.** *A matrix (linear operator) $A : M \to M$ is called **R-local** if its matrix elements $A_{xy} : M_y \to M_x$ are zero whenever the distance $d(x, y) > R$. The space of all R-local matrices is denoted $LM(S, M, R)$.*

A crucial tool for working with local operators is the following lemma, which shows that a global property (like a norm bound) can be inferred from local checks.

Lemma 4.8 (Locality Lemma)**.** *Let M be an R-local operator. If for every ball D of radius R' (with $R' \gg R$), the norm of the restriction of M to the subspace associated with D is bounded, $\|M_D\| \leq a$, then the norm of the full operator M is bounded:*

$$\|M\| \leq a\left(1 + O\left(\frac{|D|}{R'}\right)\right)$$

where the constant in the $O(\cdot)$ notation depends on the dimension n. This means that properties that hold locally on patches of a certain size extend to the entire system.

Bulk Gap for Extended Systems With the concept of locality, we can define a robust notion of a bulk energy gap that is insensitive to what happens at the boundaries. We must ignore a "margin" near the boundary of the system.

Definition 4.9 (Bulk Gap)**.** *An R-local, anti-symmetric matrix A has a* ***bulk gap*** *Δ if the matrix $-A^2$ (which is $2R$-local) is positive definite* away from the boundaries. *More precisely, for any sub-region $S' \subseteq S$ that is sufficiently far from the boundary of S, the restriction of $-A^2$ to S' satisfies:*

$$(-A^2)|_{S'} \geq \Delta^2 I$$

It is critical that we **first square the operator, then restrict** to the bulk. Squaring mixes information locally (a product of two R-local operators is $2R$-local), and only after this local information has been combined do we discard the boundary region where the physics might be gapless. This procedure can be formalized using quotient spaces, where we mod out by operators that only have support in the margin.

This local and soft framework allows for deformations like stretching or compressing parts of the lattice. So long as these deformations are smooth enough, properties like the bulk gap are preserved. This is the principle of **softness**, which, together with locality, is essential for building a robust theory of topological phases.

4.3 Lecture 3

4.3.1 Recap: Local Matrices and the Bulk Gap

We begin by recalling the framework of local matrices developed in the previous lecture.

Definition 4.10 (Space of Local Matrices)**.** *Let S be a metric space of sites, and let $S_0 \subseteq S$ be a specified boundary region. To each site $x \in S$, we associate a finite-dimensional Euclidean space M_x. The space of* ***r-local matrices*** *on the pair (S, S_0), denoted $LM(S, M, r)$, is the space of operators A on the total space $M = \bigoplus_{x \in S} M_x$ whose matrix elements $A_{xy} : M_y \to M_x$ vanish if the distance $d(x, y) > r$.*

For example, consider three sites $\{1, 2, 3\}$ on the real line, each with an associated space $\mathbb{R}^4$. A matrix element coupling sites 1 and 2 with strength J might look like:

$$A_{12} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & J & 0 \\ 0 & -J & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

where J is a real number. We will often use diagrams to represent such matrices.

The definition of a bulk energy gap was formulated to be robust against truncations of the system.

Definition 4.11 (Bulk Gap)**.** *An r-local anti-symmetric matrix A has a* ***bulk gap*** *$\Delta > 0$ if the operator $X = -A^2 - \Delta^2 I$ is positive semi-definite* after restricting to the interior of the system. *We first compute the $2r$-local matrix $-A^2$ on the full space S, and then we take its sub-matrix corresponding to the region $S \setminus margin(S_0)$, where the margin is an R-neighborhood of the boundary S_0. The matrix elements of A^2 near the boundary are considered unreliable ("garbage") due to truncation effects, and this procedure discards them.*

The framework relies on the **Locality Lemma**, which states that if a local operator has a norm bounded by a on all small balls of radius R', its global norm is also bounded, $\|X\| \leq a(1 + O(R/R'))$. This lemma is known to hold in Euclidean space and, more generally, in spaces where the volume does not grow too quickly (e.g., faster than polynomially). It breaks down in spaces with exponential volume growth, such as hyperbolic spaces or some trees.

4.3.2 Softness: Fundamental Operations

The principle of **softness** means our definitions should be robust under small, continuous deformations. This is realized through several key operations.

Wiggling and Enlarging We can always enlarge the local Hilbert spaces $M_x \to M_x \oplus \mathbb{R}^{2k}$ by adding trivial, gapped degrees of freedom (e.g., copies of the matrix $\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$). This creates more "room" in the system. This extra

room allows us to "wiggle" the system, meaning we can continuously move the lattice sites around, provided the displacement is small compared to the locality radius.

Spatial Interpolation Given a one-parameter family of local matrices $\{A(t)\}_{t\in[0,1]}$ on a space S', we can construct a single, spatially-varying matrix $\tilde{A}$ on a larger space S. Let $f : S \to [0,1]$ be a map from our physical space to the parameter space (e.g., $f(x) = x/L$).

Definition 4.12 (Spatial Interpolation)**.** *The spatially interpolated matrix $\tilde{A}$ is defined by its matrix elements:*

$$\tilde{A}_{xy} = \frac{1}{2}\left((A(f(x)))_{xy} + (A(f(y)))_{xy}\right)$$

The averaging of $f(x)$ and $f(y)$ is a technical detail to ensure $\tilde{A}$ is exactly anti-symmetric. Conceptually, the matrix at position x is simply $A(f(x))$.

This allows us to "glue" different phases together. For example, given a path $A(t)$ from a matrix A to another matrix B, we can construct a single system that smoothly transitions from phase A on one side to phase B on the other.

Folding Folding is a powerful technique, analogous to the Eilenberg-Mazur swindle, that allows us to improve the locality of a matrix at the cost of increasing the number of degrees of freedom.

The process is as follows:

1. Start with a system described by a matrix A.
2. Using spatial interpolation, create a new system composed of A and its inverse $-A$, glued at one end. This combined system $A \oplus (-A)$ can be created from a trivial system.
3. Glue the other ends of A and $-A$ together, forming a closed loop.
4. This process can be repeated. We can create any even number of layers, $A \oplus (-A) \oplus A \oplus \dots$.
5. This layered object, which exists on a 1D line, can be geometrically rearranged ("wiggled") into a new configuration where the layers are stacked vertically. This new matrix $\tilde{A}$ is homotopic to the original A.

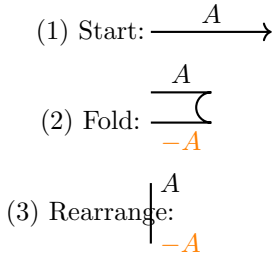

The key insight is that by stretching this new configuration in the vertical direction (which corresponds to enlarging the local Hilbert spaces M_x), we can make the matrix arbitrarily local in the horizontal (physical) direction. This allows us to take a limit where the locality radius $r \to 0$.

4.3.3 Constructing the Ω-Spectrum

These tools allow us to define a sequence of topological spaces $\{L_n\}$ and show they form an Ω-spectrum.

Definition 4.13 (Spaces of Local Gapped Matrices)**.** *Let $L_A(S, M, r, \Delta)$ be the space of r-local, anti-symmetric matrices on (S, M) with a bulk gap Δ. We define L_n to be the space of local matrices on the n-dimensional disk D^n relative to its boundary ∂D^n. We take a limit where the system size becomes infinite, the locality radius $r \to 0$, and the gap $\Delta \to 1$ (gap amplification via folding).*

$$L_n = \lim_{r\to 0, \Delta \to 1} \bigcup_M LA(D^n, M, r, \Delta)$$

While this limit is complicated, its homotopy type can be rigorously defined using concepts like persistent homology or by defining the space via its representable functor of maps from test spaces X.

To show that $\{L_n\}$ is an Ω-spectrum, we must construct homotopy equivalences $\alpha_n : L_n \to \Omega(L_{n+1})$ and $\beta_n : \Omega(L_{n+1}) \to L_n$.

The Suspension Map $\alpha_n : L_n \to \Omega(L_{n+1})$ An element of $\Omega(L_{n+1})$ is a map from the circle S^1 to L_{n+1}, i.e., a loop of $(n+1)$-dimensional local matrices based at the trivial matrix. We construct this loop by "pumping" an n-dimensional system $A \in L_n$ through the extra dimension.

The process for $\alpha_0 : L_0 \to \Omega(L_1)$:

1. Start with a trivial 1D system (a collection of decoupled, gapped dimers). Let this be the basepoint of our loop at time $t = 0$.
2. For $t \in [0, 1/2]$, we use a path that deforms a pair of trivial dimers $(A_0, -A_0)$ into our desired system and its inverse $(A, -A)$. We do this for alternating pairs along the line.
3. For $t \in [1/2, 1]$, we annihilate the pairs in a shifted pattern.
4. The net result is that after one full cycle ($t = 1$), the system returns to the trivial state, but a copy of A has been effectively transported or "pumped" from one end of the system to the other.

This process creates a loop in the space of 1D local matrices, which is an element of $\Omega(L_1)$.

The Loop Map $\beta_n : \Omega(L_{n+1}) \to L_n$ The map β_n is simply spatial interpolation. An element of $\Omega(L_{n+1})$ is a family of matrices $A(t)$ parameterized by the

loop coordinate $t \in S^1$. We construct a single n-dimensional matrix by identifying the spatial coordinate x with the loop parameter t. We slice the loop diagram diagonally.

Theorem 4.14. *The maps α_n and β_n are homotopy inverses:*

$$\beta_n \circ \alpha_n \simeq id_{L_n} \quad \text{and} \quad \alpha_n \circ \beta_n \simeq id_{\Omega(L_{n+1})}$$

Therefore, the sequence of spaces $\{L_n\}$ forms an Ω-spectrum.

The proof relies on the folding trick. These arguments are completely general and rely only on the properties of invertibility, locality, and softness. They can be applied to other systems, such as interacting Hamiltonians or quantum cellular automata, provided these properties can be established.

4.3.4 Towards Computing the Spectrum

This general construction proves the existence of an Ω-spectrum but does not tell us what it is. To compute the homotopy groups $\pi_k(L_n)$, we need more structure.

The key is to consider matrices with additional Clifford symmetries. Let $L^n_{p,q}$ be the space of n-dimensional local matrices that anti-commute with p Clifford generators ($C_i^2 = 1$) and q anti-Clifford generators ($C_j^2 = -1$).

- Real Bott periodicity for Clifford algebras implies that the algebraic structure only depends on $p - q \pmod 8$.
- The construction of the Ω-spectrum can be shown to respect these symmetries, leading to a proof of **topological Bott periodicity**: $\pi_k(L^n_{p,q}) \cong \pi_k(L^n_{p-q,0})$.

This structure provides a concrete algorithm for classifying a given free fermion system:

1. An n-dimensional system in L_n can be mapped to an equivalent $(n-1)$-dimensional system with an extra Clifford symmetry, i.e., an element of $L^{n-1}_{1,0}$.
2. By iterating this n times, any system in L_n can be mapped to an equivalent system in $L^0_{n,0}$.
3. A system in $L^0_{n,0}$ is just a single finite matrix with n Clifford symmetries. Its classifying invariant can be computed using standard linear algebra (e.g., by computing the sign of a Pfaffian).

This algorithm allows one to take a finite patch of any free fermion system and compute its topological invariant, determining which phase it belongs to in the KO-theory classification.

4.4 Lecture 4

4.4.1 Computing the Free Fermion Spectrum via Clifford Symmetries

In the previous lecture, we established that the spaces of local, gapped, anti-symmetric matrices, denoted $\{LA_n\}$, form an Ω-spectrum via the maps $\alpha_n : LA_n \to \Omega(LA_{n+1})$ and $\beta_n : \Omega(LA_{n+1}) \to LA_n$. This abstract fact does not, however, identify the spectrum. To compute it, we introduce additional structure: Clifford symmetries.

Definition 4.15 (Matrices with Clifford Symmetries)**.** *Let $Cl_{p,q}$ be the real Clifford algebra with p generators $\{e_i^+\}$ squaring to $+1$ and q generators $\{e_j^-\}$ squaring to -1. We consider a space of matrices C (e.g., symmetric or anti-symmetric) that act on a representation space of $Cl_{p,q}$ and anti-commute with all $p+q$ generators:*

$$Ce_k = -e_k C \quad \text{for all generators } e_k$$

The space of such n-dimensional, local, anti-symmetric matrices is denoted $LA_n^{p,q}$.

There is a fundamental relationship, a form of Morita equivalence, between matrices of different symmetry types.

Proposition 4.16 (Morita Equivalence)**.** *The space of anti-symmetric matrices with (p,q) Clifford symmetries is isomorphic to the space of symmetric matrices with $(q+2,p)$ Clifford symmetries.*

$$LA_n^{p,q} \cong LS_n^{q+2,p}$$

Remark 4.17 (Examples of Clifford Symmetries)**.** *A common representation for the first two positive generators on a space $M \oplus M$ is $e_1^+ = \begin{pmatrix} I & 0 \\ 0 & -I \end{pmatrix}$ and $e_2^+ = \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix}$. A symmetric matrix $C \in LS_0^{1,0}$ (anti-commuting with e_1^+) must have the block form $C = \begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix}$. If we further demand $C^2 = I$, then A must be an orthogonal matrix. This establishes a correspondence between $LS_0^{1,0}$ and orthogonal matrices.*

A key fact of Clifford algebras is Bott periodicity: the algebraic structure depends only on $p-q \pmod 8$. We will show that this algebraic periodicity manifests as a topological periodicity in our spectrum, $LA_n^{p,q} \cong LA_n^{p-q \pmod 8,0}$.

Decomposition of the Suspension Map The suspension map $\alpha_n : LA_n^{p,q} \to \Omega(LA_{n+1}^{p,q})$ can be decomposed into two more fundamental maps, μ and ν, when Clifford symmetries are present. Assume $p > 0$.

$$\alpha_n : LA_n^{p,q} \xrightarrow{\mu} LA_{n+1}^{p-1,q} \xrightarrow{\nu} \Omega(LA_{n+1}^{p,q})$$

1. **The map μ:** This map takes an n-dimensional system $A \in LA_n^{p,q}$ and produces an $(n+1)$-dimensional system. The matrix A anti-commuting with e_1^+ has a block form $A = \begin{pmatrix} B & C \\ -C^T & D \end{pmatrix}$. From these blocks, we construct a new 1D matrix—a generalized Majorana chain—with on-site and hopping terms determined by B, C, D. This new system lives in $n+1$ dimensions and has lost the e_1^+ symmetry, so it belongs to $LA_{n+1}^{p-1,q}$. The explicit construction for $\mu(A)$ is a periodic 1D chain built from the following template:

$$\mu(A) = \begin{pmatrix} I & 0 \\ 0 & -I \end{pmatrix} \otimes B + \begin{pmatrix} 0 & -I \\ I & 0 \end{pmatrix} \otimes C + \text{derivative term} \propto \begin{pmatrix} 0 & I \\ -I & 0 \end{pmatrix}$$

   This construction can be shown to preserve the energy gap.

2. **The map ν:** This is a standard construction in K-theory. To get from $LA_{n+1}^{p-1,q}$ to $\Omega(LA_{n+1}^{p,q})$, we first "double" a system $\tilde{A} \in LA_{n+1}^{p-1,q}$ to $\begin{pmatrix} \tilde{A} & 0 \\ 0 & -\tilde{A} \end{pmatrix}$. This doubled system gains two new Clifford generators, e.g., $\begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix}$ and $\begin{pmatrix} 0 & -iI \\ iI & 0 \end{pmatrix}$. Let's say we gain a negative generator e^-. We then construct a path of matrices for $t \in [0,1]$:

$$A''(t) = \cos\left(\frac{\pi t}{2}\right)\tilde{A} + \sin\left(\frac{\pi t}{2}\right)e^-$$

   This path connects $\tilde{A}$ to e^-. By creating a path from a basepoint to $\tilde{A}$, then along this path, and back, we construct a loop, which is an element of $\Omega(LA_{n+1}^{p,q})$.

Algorithmic Dimensional Reduction The power of this decomposition lies in its inverse. The composition $\beta_n \circ \nu$ gives a map that reduces dimension by one while adding a Clifford symmetry. By iterating this procedure n times, we can map any system in LA_n to an equivalent system in LA_0 with n extra Clifford symmetries. The classification problem for an n-dimensional field theory is thus reduced to a zero-dimensional linear algebra problem.

4.4.2 Beyond Free Fermions: Interacting Systems

The framework of local, invertible systems is more general than just free fermions. We can consider interacting fermionic and bosonic systems. The hierarchy of phases is:

$$LA_n(\text{free}) \subseteq F_n(\text{invertible fermionic}) \subseteq Topo_n(\text{topological}) \subseteq Gapped_n$$

We denote the corresponding spectra as LA, F, and B. Their stable homotopy groups $\pi_k(E)$ classify phases in dimension $-k$. The table below summarizes the

known and conjectured homotopy groups.[1]

n	'$\pi_n(LA) + \pi_n(F) =$'	'$\pi_n(B)$'
3	$\frac{1}{2}$	
2	$\mathbb{Z}$	$\mathbb{Z}$
1	0	0
0	$\mathbb{Z}_2$	$\mathbb{Z}_2$
-1	$\mathbb{Z}_2$	$\mathbb{Z}_2$
n'		
0	0	0
1	$\mathbb{Z}$	$\mathbb{Z}$
2	$\mathbb{Z}_2$	$\mathbb{Z}_2$

A remarkable relation exists between the bosonic and fermionic worlds. The generator of $\pi_{-2}(B) \cong \mathbb{Z}$ is the E8 state, a 2D bosonic phase whose edge theory is the E_8 WZW model ($c_- = 8$). The generator of $\pi_{-2}(LA) \cong \mathbb{Z}$ is the p+ip superconductor ($c_- = 1/2$). The inclusion map $U : LA \to F$ and the forgetful map $V : B \to F$ relate these generators:

$$16 \cdot U(\text{p+ip}) = V(E_8)$$

4.4.3 Connection to Cohomology and Physics

The structure of these spectra can be understood through the lens of generalized cohomology theories.

- The space of states in 0 dimensions, B_0, is the space of normalized vectors in a Hilbert space, which is CP^∞. Thus, $B_0 \simeq CP^\infty \simeq K(\mathbb{Z}, 2)$.
- This implies the existence of a map from the Eilenberg-MacLane spectrum to the spectrum of bosonic phases: $I\mathbb{Z}[2] \to B$.
- This abstract map has a physical realization. Given a symmetry group G, it produces a map from group cohomology to SPT phases: $H^{n+2}(BG; \mathbb{Z}) \to [BG, B_n]$. This provides a top-down derivation for the classification of Symmetry Protected Topological phases.

A map in the other direction, from phases to cohomology, $B_n(X) \to I\mathbb{Z}^{n+2}(X)$, also exists. It can be constructed abstractly from the structure of the spectra. Physically, this map is realized by the **higher Berry curvature**.

This corresponds to a physical observable: the **chiral central charge** c_-. For a gapped 2D system, the edge will have a thermal energy flow given by:

$$\text{Energy flow} = \frac{\pi c_-}{12\hbar} T^2$$

[1] The table is transcribed literally from the blackboard. The column headers are unconventional. The column '$\pi_n(LA) + \pi_n(F) =$' seems to list values for $\pi_n(F)$, which are expected to match $\pi_n(LA)$ where defined. Standard K-theory gives $\pi_{-3}(LA) = \pi_{-1}(KO) = \mathbb{Z}_2$, not $\frac{1}{2}$. The entry '1/2' may refer to a specific physical quantity like a chiral central charge.

The map $B_2 \to H^4(X, \mathbb{Z})$ evaluated on a phase with chiral central charge c_- gives an integer value of $2c_-$. The argument implies that c_- is well-defined modulo 12. Further results by Nikita Satienko show it is well-defined modulo 24. While it is believed that c_- is a well-defined rational number for any gapped phase, a complete proof is still missing. The existence and quantization of this invariant is crucial for establishing that the maps between the known phases (LA, F, B) are isomorphisms.