BOTT PERIODICITY AND THE "PERIODIC TABLE" OF TOPOLOGICAL INSULATORS AND SUPERCONDUCTORS

by

Martin R. Zirnbauer

Abstract. **—** Bott periodicity is said to be one of the most surprising phenomena in topology. Perhaps even more surprising is its recent appearance in condensed matter physics. Building on work of Schnyder et al., Kitaev argued that symmetry-protected ground states of gapped free-fermion systems, also known as topological insulators and superconductors, organize into a kind of periodic table governed by a variant of the Bott periodicity theorem. In this colloquium, I will sketch the mathematical background, the physical context, and some new results of this ongoing story of mathematical physics. —

The text below is that of a one-hour Mathematical Colloquium, delivered at LMU Munich (Oct 2014), IAS Dublin (Nov 2014), HIM Bonn (Dec 2014), SCGP Stony Brook (Sept 2015), Harvard (Oct 2015), Geneva (Journée de Rham, March 2016), Erlangen (May 2016), Bristol (Nov 2016), Cologne (April 2017), Aachen (April 2017), Paderborn (July 2017).

1. Overture

The story of this colloquium is centered in condensed matter theory, but it has the very attractive feature of great scope, ranging from the experimental physics of topological insulators to the pure mathematics of algebraic topology. This makes it one of the best stories, in my opinion, in theoretical physics of the last decade.

The plan of the colloquium is to tell the story in 3 parts: first some impression of the relevant physics, then some mathematical background, and finally how the two come together.

2. Physics background

2.1. Bott-Kitaev Table. — Here [diagram] is a wonderful principle due mainly to Kitaev, following work by the UCSB group: the "Periodic Table" (of topological insulators and superconductors), which caught the attention and fancy of condensed matter physics a few years ago and forms the theme of this colloquium.

Let's begin with some history and context.

The intense research activity on topological insulators started about 10 years ago after the theoretical and experimental discovery of the Quantum Spin Hall Insulator. As the name suggests, this is a close cousin of the standard Quantum Hall Effect, although it differs from it by the presence of spin and by time-reversal symmetry (meaning invariance under the hypothetical operation of inverting the time direction).

By 2008 about a half a dozen of related systems had been identified; these included superfluid ^{3}He in the B phase, the Majorana chain, and more. It was then Kitaev who realized that there exists some organizational principle (related to Bott periodicity), which ties all these examples together, and he assembled them into a kind of Periodic Table (reprinted here from a Review of Modern Physics). The entries of the table are meant to be sets of equivalence classes of ground states of many electrons (treated in mean-field approximation) as a function of space dimension d and a variable number of certain symmetries. (And space here means momentum space approximated as a sphere).

For me, this Table posed a particular challenge (if not provocation) as my name was and is often cited in this context. The Table took me initially by surprise, for my work with Altland of the mid-nineties concerned the symmetry classes of disordered electrons and their relation to the Cartan classification of symmetric spaces, whereas the subject matter here is something quite different, namely symmetry-protected topological ground states of electron systems with an energy gap for excitations.

2.2. IQHE. — Now I suspect that most of my mathematical audience is not familiar with the physics notion of topological insulators. I will therefore say a few words of introduction, highlighting the main examples. (Please be warned that this will be impressionistic, as I don't know how to explain physics to mathematicians, not in 15 minutes anyway.)

Paying tribute to history (old and recent), our first example is the Integer Quantum Hall Effect; which takes place in a two-dimensional disordered electron gas in a very strong magnetic field and at very low temperatures. The experimental observations mean that in certain ranges of the magnetic field a voltage gradient does not drive a flow of electrical current in the direction of the gradient; thus the electron gas is an electrical insulator. However, this insulator has the striking feature (indicating its topological nature) that there exists current flow along the lines of constant voltage. Thus a current I along the sample comes with a voltage drop V_H across the sample. The constant of proportionality, the Hall conductance I/V_H, is measured to be an integer n (with astounding accuracy) in units of the conductance quantum e^2/h.

By an insight of David Thouless (who received part of the Physics Nobel Prize 2016 for it) this integer n has an interpretation as the Chern number of the complex line bundle modeling the ground state of the electron gas. More precisely, the topological invariant of Chern number characterizes the ground state of the bulk, i.e. the infinite system without boundary. In reality, there always exists a boundary, and then the topological invariant manifests itself as n conducting edge channels. The heuristic principle behind this duality is called the bulk-boundary correspondence.

2.3. QSHI. — Our second example is the Quantum Spin Hall Insulator. This topological state was first observed in a layer of HgTe. More generally, it occurs in 2d materials with strong crystal-electric fields and electrons fast enough in order for the relativistic phenomenon of **spin-orbit coupling** to become noticeable. The role of strong spin-orbit coupling is to cause band inversion and thereby put a topological "twist" in the ground state of the electron gas. More precisely, ground states of the type at hand (i.e., in two dimensions, with time-reversal invariance and conserved charge) are classified by a topological invariant called the Kane-Mele index, and this invariant may become non-trivial when the spin-orbit coupling is strong.

By the principle of bulk-boundary correspondence, a non-trivial Kane-Mele index implies a pair of counter-propagating edge modes, thus endowing the 2d insulator with a 1d surface that is actually **conducting**. The beauty of this effect is that it is robust with respect to the addition of disorder and weak electron-electron interactions, and yet it already occurs and can be studied in the simple setting of independent electrons with translation invariance.

And last but not least, the conducting surface states have been seen in transport experiments (showing a conductance of 2 in the natural units of conductance quantum).

2.4. Majorana chain. — To give a third example (of a topological ground state) of great experimental interest, let me mention the so-called Majorana chain (a.k.a. Kitaev chain). For a theorist, this is a one-dimensional superconductor of spinless electrons with no symmetries at all. In reality, of course, the superconductor is *quasi*-one-dimensional and one has to do some material science and trickery to make spinless electrons out of real electrons with spin, break time reversal invariance, and still induce superconductivity (to break the symmetry of charge conservation).

This example is simple enough for the topological twisting to be explained in a few formulas and a picture. Assuming that momentum is conserved and that momentum space is a circle parameterized by $k \in \mathbb{R}/2\pi\mathbb{Z}$, let

$$\gamma_k = u(k)\, c_k + v(k)\, c_{-k}^\dagger$$

be the transformation from the bare electron annihilation operator c_k to the quasi-electron annihilation operator γ_k of the superconducting ground state. Note that γ_k determines a complex line in $\mathbb{C}^2$, hence a point in $\mathbb{C}P^1$, for every value of the momentum k.

Now Fermi statistics dictates that all quasi-electron annihilation operators must anti-commute with one another:

$$\{\gamma_k, \gamma_{-k}\} = 0\,.$$

This leads to a condition

$$u(k)\, v(-k) + v(k)\, u(-k) = 0\,,$$

which I call the *Fermi constraint*. Its solutions of interest (from a topological perspective) are those with

$$u(0) = v(\pi) = 0\,.$$

Thus the quasi-electron annihilation operator for large momenta $k = \pm\pi$ is the bare annihilation operator, and for zero momentum it is the bare *creation* operator, drawn in [this picture] as a point and its antipode in $\mathbb{C}P^1$. What is crucial here is that this feature is *rigid* (not allowing deformation) – you can deform the intermediate part of the curve, but the ends are nailed down by the Fermi constraint. This rigidity results in a topological invariant, which actually makes the superconducting ground state (of the Majorana chain) topologically non-trivial.

By bulk-boundary correspondence, a finite Majorana chain then has zero modes at its ends (so-called Majorana end states), which are topologically protected and hence robust with respect to disorder, interactions, and in fact just about anything. There exist dozens and dozens of proposals how this striking effect of protected zero modes might be used for the purpose of building a fault-tolerant quantum computer ...

I'll leave it at that, putting on record that QSHI, the Majorana chain and, more generally, the field of topological insulators and superconductors, is a field with an extreme amount of activity, spurred by its mathematical fascination and a great potential for experimental applications.

Now, if you are a mathematician and didn't understand anything, I apologize. As a small compensation, here is an alternative definition of the notion of superconductor [cartoon].

A one-line summary is that the ground state of a topological insulator or superconductor carries a topological invariant which sets it apart from a trivial insulator and makes it physically very interesting, mostly by its conducting surface states.

3. Bott Periodicity

Let me now continue with mathematical background to give some meaning to the first word in my title: *Bott periodicity*.

3.1. Front page of Bott's 1959 paper. — Let X be a compact Riemannian manifold, and fix a pair of points p and q together with a homology class h of paths connecting them. Abbreviate this data set as $\nu \equiv (p,q;h)$. Let X^ν denote the space of minimal geodesics (of class h) joining p with q. (Of course, this space is boring for q close to p, but it becomes interesting when q is a conjugate point of p.) Now recall that the geodesics of a Riemannian manifold are extrema of the path length function. Consequently, one can associate with every geodesic a Hessian. Define the index, $|\nu|$, of ν as the minimum number of negative eigenvalues (or unstable modes) of this Hessian, where the minimum is taken over the set of non-minimal geodesics of type ν, i.e. geodesics going from p to q in class h.

3.1.1. Bott's Theorem.— In this setting, Bott proved the following beautiful result. If the Riemannian manifold X is a *symmetric* space, then the space (of minimal geodesics) X^ν is another symmetric space, and the latter furnishes a good approximation to the loop space of the former in the sense that the homotopy groups agree in low dimension (up to a threshold set by the index of ν):

$$\pi_d(X^\nu) = \pi_{d+1}(X) \quad \text{for} \quad 0 < d < |\nu| - 1 .$$

3.1.2. Example. — To illustrate, take X to be the n-sphere and consider a pair of antipodal points, say the North Pole and the South Pole. (There is no homology class to be specified here, as the sphere is simply connected.) The minimal geodesics from the North Pole to the South Pole are parameterized by the equator, which is an $(n-1)$-sphere:

$$X = \mathrm{S}^n, \quad X^{(p,-p)} = \mathrm{S}^{n-1}, \quad |\nu| = 2(n-1).$$

By computing the associated index and applying his Theorem to this particular situation, Bott was able to reproduce a result already known to Freudenthal:

$$\pi_d(\mathrm{S}^{n-1}) = \pi_{d+1}(\mathrm{S}^n) \quad \text{for} \quad 0 < d < 2n-3 .$$

3.2. What's a symmetric space?— The statement of Bott's Theorem assumes that you are familiar with the notion of symmetric space. Just in case you are not — I guess there might be a few students in the audience — here are a few words of explanation:

In Riemannian geometry there exists something called the Riemann curvature tensor. In a coordinate basis, it has the well-known expression

$$R^i{}_{jkl} = \partial_k \Gamma^i_{lj} - \partial_l \Gamma^i_{kj} + \Gamma^m_{lj}\Gamma^i_{km} - \Gamma^m_{kj}\Gamma^i_{lm} .$$

A (locally) symmetric space is defined to be a Riemannian manifold $X = U/K$ with a Riemann tensor which is covariantly constant:

$$\nabla R = 0 .$$

The simplest example of a such a space is the round two-sphere S^2,

$$X = \mathrm{S}^2, \quad ds^2 = d\theta^2 + \sin^2\theta\, d\phi^2 ,$$

with line element ds^2 induced by the Euclidean distance of three-dimensional space. A more general example is the Grassmann manifold

$$X = \mathrm{Gr}_n(\mathbb{C}^N) = \mathrm{U}(N)/\mathrm{U}(n)\times\mathrm{U}(N-n),$$

of complex n-planes in $\mathbb{C}^N$, which may also be viewed as the quotient of the unitary group (by a suitable subgroup). In the case of $n = 1$ and $N = 2$ this is also known as $\mathbb{C}P^1$, which we met in the example of the Majorana chain.

An important and well-known fact about symmetric spaces is that they were completely classified [the globally symmetric ones] by the French geometer Elie Cartan. Aside from a finite number of exceptional spaces, they come in 10 large families, which Cartan called A, AI, AII, AIII, BD, BDI, C, CI, CII, and DIII. (Please take note of the number 10.)

3.3. Complex and real Bott ν-sequence. — Returning to our subject of Bott periodicity, we make the observation that the process of forming X^ν from X can be iterated to produce a *sequence* of symmetric spaces. Bott speaks of a ν-*sequence*.

Here [diagram] are the two ν-sequences of relevance to our subject; they are called complex and real for short. The complex one stems from the fact that the minimal geodesics (connecting a point with its antipode) in a complex Grassmann manifold are parameterized by points in a unitary group, and vice versa. This sequence has period 2.

The real ν-sequence is more involved and has period 8. One may start with a symmetric space O/U, which you may think of as the space of complex structures of an even-dimensional real vector space. The minimal geodesics between antipodes of O/U turn out to be in bijection with points of U/Sp, whose minimal geodesics in turn are in bijection with the points of a symplectic Grassmann manifold, and so on.

3.4. Morse Theory. — Let me now embark on a mathematical digression to sketch the reasoning by which Bott arrived at his famous theorem. For this purpose, let's suppose that we want to determine the homotopy type of some topological space M. This can be done by using Morse theory to construct a cell decomposition for M, as follows.

One picks a "good" function $f : M \to \mathbb{R}$ and looks at the so-called half-spaces M^a for f; the boundaries of these half-spaces are the level sets of f. (For example, M might be a 2-torus, f could be the height function, and M^a would then be the part of the torus with height no greater than a.)

The first statement of Morse theory is this: if f has no critical value between a and b, then the half-spaces M^a and M^b have the same *homotopy type*; or in non-technical language: they are topologically equivalent. And, indeed, if we fix some intermediate height a in our torus example, then M^a is a cylinder, and remains so under any small variation of a.

The second statement of Morse theory is more involved: let $x \in M$ now be a critical point (or $f(x)$ a critical value) and let the Hessian of f at x have n negative eigenvalues and *no* zero eigenvalue; assume further that x is the *only* critical point for some ε-neighborhood of the critical value $f(x)$. Then, when we pass through the critical value $f(x)$ the homotopy type changes in a specific way, namely by attaching an n-cell (or n-disk):

$$M^{f(x)+\varepsilon} \sim M^{f(x)-\varepsilon} \cup e_n .$$

To illustrate: in our example, if we pass through the critical value c_1, then the homotopy type changes from that of a disk (which contracts to a point) to that of a cylinder (which contracts to a circle), and this change amounts to attaching an n-cell where $n = 1$ is the number of unstable modes at the saddle point of height c_1.

I am trying to give you a glimpse of a some mathematics which is not all that simple ... The upshot is an iterative scheme: take f to be a Morse function and sweep through its half-spaces from the empty set to the full space by raising the level. Start from nothing (for the empty set). At every encounter of a critical value of index n attach an n-cell. By this constructive process you build the skeleton of a cell decomposition, which encodes all the topological information about M.

3.5. Idea of Bott's proof. — How did Bott use this machinery to prove his famous theorem? Here it is:

Given the Riemannian manifold X and a triple $\nu = (p,q;h)$, take for M the space of paths from p to q of homology class h (and denote this by $\Omega^\nu X$). Using the Riemannian geometry of X, introduce the function f (a Morse function) which assigns to every path its length. Then, for any length a smaller than the geodesic distance L between p and q, the half-space M^a is empty. Now the critical points of the length function f are geodesics of the Riemannian manifold X, and if you increase the length parameter a a little bit beyond L, the half-space turns into the space X^ν of minimal geodesics:

$$M^a = X^\nu \quad \text{if} \quad L < a < L + \varepsilon\,,$$

Increasing the length further, you eventually encounter non-minimal geodesics; their existence makes itself felt by n-cells that need to be attached to X^ν:

$$\Omega^\nu X = X^\nu \cup e_n \cup e_n' \cup \ldots \quad \text{where} \quad \dim e_n \geq |\nu|\,.$$

Now, the dimension of any such cell is bounded from below by $|\nu|$, which we had defined as the minimal number of unstable modes for any non-minimal geodesic. If this number $|\nu|$ is large in comparison with the dimension d, then (by another fundamental result of Morse theory) the cells attached to X^ν don't matter and the homotopy groups of the path space are the same as those of X^ν (which in turn are the same as the homotopy groups of X in dimension $d+1$):

$$\pi_{d+1}(X) = \pi_d(\Omega^\nu X) = \pi_d(X^\nu) \quad \text{if} \quad 0 < d < |\nu| - 1.$$

3.6. Clifford algebra formulation. — So far, I have told two very different stories (Bott periodicity and topological insulators) and you might be wondering how they are going to be related. Well, the current slide is a good preparation for that.

Consider a complex vector space W of dimension $2n$, and let W carry two structures: a Hermitian scalar product and a (non-degenerate) symmetric bilinear form (or bracket)

$$\{\cdot,\cdot\} : W \otimes W \to \mathbb{C}.$$

Assume that we are given a system of operators $J_1, \ldots, J_s$ on W which are unitary with respect to the Hermitian structure, orthogonal with respect to the bracket structure, and satisfy the relations of a Clifford algebra represented on W:

$$J_l J_m + J_m J_l = -2\delta_{lm} \mathrm{Id}_W\,, \quad (J_l^* = J_l^{-1} = J_l^T).$$

Observe that any such operator J has only two eigenvalues, $+\mathrm{i}$ and $-\mathrm{i}$, and that we may decompose W by the eigenspaces associated with these two eigenvalues.

Now let $C_0(n)$ be the Grassmann manifold of complex n-planes in W and make the following definition:

$$\begin{aligned} C_s(n) &:= \{A \in C_0(n) \mid J_1 A = \ldots = J_s A = A^{\mathrm{c}}\}, \\ R_s(n) &:= \{A \in C_s(n) \mid \{A,A\} = 0\}. \end{aligned}$$

Thus $C_s(n)$ is the subspace of complex n-planes A that are sent to their orthogonal complement A^{c} by each of the Clifford generators $J_1, \ldots J_s$. And inside of $C_s(n)$ we have $R_s(n)$, the subspace of A's that also have the property of being isotropic w.r.t. the bracket of W.

Then there is the following statement. The sequence $C_0(n) \supset C_1(n) \supset C_2(n) \supset \ldots$ realizes the *complex* Bott sequence, alternating between the complex Grassmannian and the unitary group, while $R_0(n) \supset R_1(n) \supset R_2(n) \supset \ldots$ realizes the *real* Bott sequence, which has period 8. This is spelled out in [Table], where one can also see the occasional effect of homology.

3.6.1. Sketch of proof for $R_0(n)$. — This result is well-known and is explained nicely in the last chapter of Milnor's book on Morse theory. Let's take a quick look at the case of $R_0(n)$. For this, let A be isotropic, with orthogonal complement A^{c}, and consider its image under a linear transformation g. If g is unitary, then we still have the decomposition $W = gA \oplus g \cdot A^{\mathrm{c}}$, and if g is orthogonal, then the vanishing of the bracket on A is also preserved. Thus the intersection of $\mathrm{U}(W)$ and $\mathrm{O}(W)$, which is a compact real form of the orthogonal group, acts on $R_0(n)$. This O_{2n}-action is easily shown to be transitive. The isotropy group of any A is a unitary group $\mathrm{U}(A) \cong \mathrm{U}_n \subset \mathrm{O}_{2n}$. Hence $R_0(n) \cong \mathrm{O}_{2n}/\mathrm{U}_n$, the initial space of the real Bott sequence, as claimed.

3.7. Bott Map. — In this setting with Clifford generators, the Bott construction of minimal geodesics can be described very concretely. Let's look at the real case, making the identifications $X \equiv R_{s-1}(n)$ and $X^{\nu} \equiv R_s(n)$. For the initial point of the paths in $\Omega^{\nu} X$ we take $E_{+\mathrm{i}}(J_s)$, the eigenspace for eigenvalue $+\mathrm{i}$ of the last generator (J_s), and for the final point we take $E_{-\mathrm{i}}(J_s)$. Given $A \in R_s(n)$, we act on it by the one-parameter group which is obtained by exponentiating the product of J_s with the J that has A and A^{c} for its $\pm\mathrm{i}$-eigenspaces. The outcome is a minimal geodesic connecting $E_{+\mathrm{i}}(J_s)$ with $E_{-\mathrm{i}}(J_s)$ by passing through A — the drawing should remind you of a picture I showed earlier.

This is the Bott Map, assigning to a point A in X^{ν} a minimal geodesic in X.

Let me mention that the Bott periodicity theorem, today, is usually approached from the algebraic perspective of K theory and its generalizations. Yet, Bott's original statement and proof has a very geometric flavor, as I have tried to indicate.

3.7.1. Bott map understood. — Returning to the complex and real ν-sequences (of before), and applying Bott's Theorem (from Morse Theory) we see that the homotopy groups stay the same if one takes one step up in space dimension and one step back in the ν-sequence. For the reasons indicated, the correspondence breaks down when the space dimension d becomes too large as compared with the relevant index. In the so-called stable regime, i.e. before the breakdown occurs, all homotopy groups in the complex and real table are determined by the fundamental groups given in the first column.

The message to take away from here is that there exists a mapping – called the Bott Map for short – whereby the points of a symmetric space are sent to minimal geodesics in a neighboring symmetric space, and this correspondence is a homotopy equivalence in low dimension.

4. Free-Fermion Ground States

After this exposition of foundational mathematics, let's return to physics and explain how Bott periodicity is reflected in the ground states of topological insulators and superconductors, assuming the so-called the free-fermion limit.

4.1. Setting. — First of all, let's fix some language: in the following, the word quasi-particle vacuum is synonymous with free-fermion ground state, which is synonymous with Hartree-Fock-Bogoliubov mean-field ground state (i.e. our electrons do interact, but we treat interactions in the most general kind of mean-field approximation).

Using the formalism of second quantization, we introduce Fock operators annihilating or creating a single particle, as before: c_α ($c_\alpha^\dagger$) . We take U to be the complex vector space spanned by the *annihilation* operators — V is the same thing for particle *creation* operators:

$$U = \mathrm{span}_{\mathbb{C}}\{c_1, \ldots, c_n\}, \quad V = \mathrm{span}_{\mathbb{C}}\{c_1^\dagger, \ldots, c_n^\dagger\}.$$

(For simplicity of discussion, I take the dimension of Hilbert space to be finite.) The direct sum of U and V then is the space of all fermion field operators:

$$W = U \oplus V \simeq \mathbb{C}^{2n}.$$

We sometimes refer to it as *Nambu space*.

Nambu space comes with some canonical structure. Most importantly, there is a non-degenerate symmetric bilinear form on W by the canonical anti-commutation relations (CAR) for Fock operators:

$$\{\cdot,\cdot\}: \quad W \otimes W \to \mathbb{C}.$$

Second, Hermitian conjugation in Fock space induces a complex antilinear involution

$$\gamma: \ W \to W.$$

Note that $\{\cdot,\cdot\}$ together with γ determines a Hermitian structure on W; thus W is exactly the same thing as the W in the Clifford-algebra setting for Bott periodicity.

Now it is a basic fact of many-body theory that a free-fermion ground state $|\text{g.s.}\rangle$ is uniquely determined by specifying the quasi-particle operators $\widetilde{c}_1, \ldots, \widetilde{c}_n$ that annihilate it,

$$\widetilde{c}_\alpha |\text{g.s.}\rangle = 0 \quad (\alpha = 1, \ldots, n).$$

Put differently, such ground states are in one-to-one correspondence with polarizations,

$$|\text{g.s.}\rangle \overset{1:1}{\leftrightarrow} W = A \oplus A^{\mathrm{c}},$$

where the complex vector space $A \subset W$ is spanned by the quasi-particle annihilation operators

$$A = \mathrm{span}_{\mathbb{C}}\{\widetilde{c}_1, \ldots, \widetilde{c}_n\} \subset W,$$

and its orthogonal complement A^{c} by the quasi-particle creation operators.

Now we observe that even in the complete absence of all symmetries, our vector spaces A are constrained by the condition that any pair of annihilation operators must have vanishing anti-commutator by the canonical anti-commutation relations. Thus the CAR bilinear form restricted to A is the zero form:

$$\{A, A\} = 0.$$

We refer to this condition as the Fermi constraint.

This is the basic structure (akin to the earlier setting for Bott periodicity), which we now wish to refine with an algebra of *symmetries*.

4.2. What's a symmetry?— To begin, we recall what is meant by a symmetry in the context of quantum mechanics. Following Wigner, this is primarily a transformation on the rays of Hilbert space, with the property that all transition probabilities are preserved.

Now one corner stone of our subject is a theorem (attributed to Wigner) stating that any quantum mechanical symmetry T lifts to Hilbert space as a linear operator, say $\hat{T}$, which is *either unitary, or anti-unitary*. In the former case the Hermitian scalar product of Hilbert space is preserved, in the latter case it is preserved up to complex conjugation.

Two remarks are in order here. First comes the obvious statement that symmetries always form a group, G. Indeed, if two operators are symmetries, then so is their composition.

The second remark is that in order for an operator $\hat{T}$ to be a symmetry of a quantum system with Hamiltonian H, we require that $\hat{T}$ *commutes* with H.

4.3. Threefold Way. — The subject of symmetries and their consequence for the classification of fermion systems (particularly with disorder) is very dear to me and I have given many a colloquium on it. Today, I must be extremely brief, or else I'll never get to the punch line of this colloquium. So, let me just mention that there exists something called the "Threefold Way" by Freeman Dyson, which classifies random-matrix ensembles based on unitary and anti-unitary symmetries in the sense of Wigner.

4.4. Tenfold Way. — By extending Dyson's scheme to Nambu space, we established (with Heinzner and Huckleberry) a symmetry classification (the "Tenfold Way") of disordered fermion systems in the free-fermion approximation (including superconductors). In short: we showed that there exist 10 symmetry classes, and these are in one-to-one correspondence with the 10 large families of symmetric spaces.

4.5. Symmetries of the Tenfold Way. — Now, in order for the sequel to make sense, I must spend a quick word on the symmetry operations of the Tenfold Way. There exist two distinguished anti-unitary symmetries. First, time reversal, which preserves the decomposition of field space into creation and annihilation operators, and squares to $+1$ (-1) for integer spin (resp. half-integer spin). And, second, particle-hole conjugation, which exchanges U and V (or particles and holes), while squaring to the identity.

On the unitary side, anything is allowed, but a particular focus is on two symmetries. First, the $\mathrm{U}(1)$-symmetry of charge or particle-number conservation, which transforms creation operators by some phase and annihilation operators by the opposite phase. Second, the elements of the spin-rotation group, $\mathrm{SU}(2)$ or $\mathrm{Spin}(3)$, generated by operators S_1, S_2, S_3.

4.6. Universal model (including symmetries). — We are now ready to incorporate the symmetry operations of the Tenfold Way into the model of free-fermion ground states as polarizations. In that process, as outlined by Kitaev, true physical symmetries are converted into what I propose to call "pseudo-symmetries". Let us write down the outcome first and do the explaining afterwards.

For a system with symmetry index s one has a Clifford algebra of s pseudo-symmetries:

$$J_l J_m + J_m J_l = -2\delta_{lm}\mathrm{Id}_W\,, \quad J_l \in \mathrm{U}(W) \cap \mathrm{O}(W).$$

where the $J_1, \ldots, J_s$ are complex-linear operators on W. A concise summary of the effect of imposing symmetries is the following definition.

Definition. – *A free-fermion ground state of symmetry class s is a polarization $W = A \oplus A^{\mathrm{c}}$ where the complex vector space A is subject to*

1. Fermi constraint: $\{A\,, A\} = 0$;

2. Pseudo-symmetries: $J_1 A = \ldots = J_s A = A^{\mathrm{c}}$.

We speak of pseudo-symmetries because each J $(J_1, \ldots, J_s)$ sends A to its orthogonal complement, whereas a true unitary symmetry would map A to itself.

Recall from earlier that solving these conditions for any increasing number s of pseudo-symmetries we obtain the two Bott ν-sequences, complex and real.

Example. – Let the gapped system (and hence the ground state) be time-reversal invariant. Then, applying time reversal T to any annihilation operator we get another annihilation operator: $TA = A$. If we carry on by applying the operator γ of Hermitian conjugation $(\gamma\colon\ c \leftrightarrow c^\dagger)$, we end up with a creation operator. Thus the composition $J_1 = \gamma \circ T$ sends A to its orthogonal complement A^{c}, which yields the first pseudo-symmetry relation: $J_1 A = A^{\mathrm{c}}$. Moreover, being the product of two anti-unitary operators, J_1 is unitary; and assuming the case of fermions with half-integer spin, T squares to minus one, J_1 does the same, and we have identified the first generator J_1 of the Clifford algebra of pseudo-symmetries.

4.7. Kitaev sequence. — How does this story continue? While Kitaev wrote down only the first two steps, in the following table we present the whole answer.

class	symmetries	s	pseudo-syms.
D	none	0	Fermi constraint
DIII	T (time reversal)	1	$J_1 = \gamma T$
AII	T, Q (charge)	2	$J_2 = \mathrm{i}\gamma TQ$
CII	T, Q, C (ph-conj.)	3	$J_3 = \mathrm{i}\gamma CQ$
C	S_1, S_2, S_3 (spin rot.)	4	see paper
CI	S_1, S_2, S_3, T	5	
AI	S_1, S_2, S_3, T, Q	6	
BDI	S_1, S_2, S_3, T, Q, C	7	

Remember that even if there are no symmetries at all, we still have the Fermi constraint due to Fermi statistics; that characterizes class D, or $s = 0$. As we have seen, imposing T gives one pseudo-symmmetry J_1, which puts us in the superconducting class DIII, or $s = 1$. Next, we add the requirement that particle number (or charge Q) be conserved. This allows us to form a second pseudo-symmetry $J_2 = \mathrm{i}QJ_1$, which is readily seen to anti-commute with J_1 and square to minus the identity. We have now arrived in class AII (housing the Quantum Spin Hall Insulator). Next, to move on to class CII or $s = 3$, we include particle-hole conjugation (C) as a third symmetry – more precisely, as a true physical symmetry which commutes with the Hamiltonian. To continue even further, we wipe the plate clean by erasing all symmetries T, Q, and C, and we demand instead that the spin-rotation generators S_1, S_2, S_3 be symmetries. (Some effort is needed to prove that the true symmetries S_1, S_2, S_3 amount to the presence of 4 pseudo-symmetries; this proof is written out in complete detail in my CMP article with R. Kennedy.) The rest of the story is a repetition of what happened at the beginning.

The 8 "real" classes in view still have to be complemented by 2 "complex" ones. To realize class A, one takes charge (Q) to be conserved. This effectively cancels the Fermi constraint and one is left with no pseudo-symmetries at all. A free-fermion ground state in this case are just plain complex vector space (or a Fermi projection). Finally, by imposing an additional symmetry of particle-hole conjugation, one arrives at class AIII, featuring one pseudo-symmetry.

To summarize, by imposing physical symmetries in a certain prescribed order, and then converting these true symmetries into pseudo-symmetries, one manages to build precisely the algebraic structure of the Clifford algebra setting for Bott periodicity (real and complex)!

5. Diagonal Map

So much for that. In the final part of the talk, I wish to present something that we call the Diagonal Map – it will be seen to be a close relative of the Bott Map.

Here is again the Periodic Table of topological insulators and superconductors; we now call it the Bott-Kitaev (Periodic) Table. Hopefully, by now you have an inkling of how our two different stories of Bott periodicity and the Tenfold Way come together. In particular, we have seen that the Kitaev sequence puts symmetry classes (of ground states) in a certain fixed order, and this order is exactly the order of Bott's ν-sequence. For me, that's quite remarkable (in the Tenfold Way, as I originally understood it, there was no preferred order).

Now, the Table should strike you as looking very similar to that of Bott but for one major difference: we have constancy of the entries along the lower diagonal (as opposed to the upper diagonal of before). Thus we are facing a phenomenon different from Bott periodicity, but presumably closely related to it.

In particular, the constancy of the entries along the lower diagonal begs a question: is this constancy just some abstract nonsense (taking place at the level of K theory, say), or is there a concrete correspondence analogous to the Bott Map – let's call it the "Diagonal Map" – which takes a topological insulator in d dimensions and transforms it into a topological insulator in one extra dimension and in the neighboring symmetry class?

The answer does not come for free but requires some serious analysis. To go further, we must be more precise about the mathematical framework used. This brings us to the issue of disorder.

5.1. Clean & Disordered Systems. — In the absence of disorder, we have translation invariance, which lets us label the quasi-particle annihilation operators by their momentum $k \in M$. Thus we get a decomposition of A as an orthogonal sum (actually, integral),

$$A = \bigoplus\nolimits_{k \in M} A_k \, .$$

Now if the physical system were a metal, the annihilation space A_k as a function of k would jump at the Fermi surface; but since we are considering gapped systems (or insulators), the assignment $k \mapsto A_k$ has the good feature of being continuous everywhere, giving a vector bundle $\{A_k\}_{k \in M}$.

If disorder is present (which is usually the case in real-world systems), we no longer have a vector bundle to work with, but we can turn to the non-commutative geometry approach which was developed for the integer quantum Hall effect 20 years ago by Bellissard and coworkers. Here one works with a C^*-algebra $\mathscr{A}$ of bounded operators and the pairing between its K-group $K_0(\mathscr{A})$ and cyclic cohomology (something invented by Alain Connes). The Hall conductance is then recognized as a non-commutative Chern number. All this is well understood for the complex symmetry classes, but is still in a state of infancy and development for the real classes – here are a few references [display list].

In the sequel I will assume the clean limit and work with the vector-bundle picture.

5.2. Diagonal Map: heuristic. — To begin, let me give a quick heuristic as to what is missing (for our purposes) from the original Bott construction and needs to be added to arrive at the Kitaev Table (of Topological Insulators and Superconductors).

Under the decomposition $A = \bigoplus_{k \in M} A_k$ (conserved momentum) the Fermi constraint $\{A, A\} = 0$ refines to

$$\{A_k, A_{-k}\} = 0 \quad (\text{for all } k \in M).$$

Thus our free-fermion ground states are vector bundles $\pi : \mathscr{A} \to M$ subject to a $\mathbb{Z}_2$-equivariance condition

$$\begin{array}{ccc} \mathscr{A} & \xrightarrow{\tau_{\mathscr{A}}} & \mathscr{A} \\ \pi \downarrow & & \downarrow \pi \\ M & \xrightarrow{\tau_M} & M \end{array}$$

with non-trivial involution $\tau_M : M \to M,\ k \mapsto -k$ and $\tau_{\mathscr{A}} : \mathscr{A} \to \mathscr{A},\ A_k \mapsto A_k^{\perp}$.

(**Warning:** $\tau_{\mathscr{A}}$ not defined on individual vectors. Thus $\mathscr{A}$, in general, is not a Real vector bundle in the sense of Atiyah.)

5.3. Diagonal Map $(d,s) \to (d+1,s+1)$**. —** The goal now is to take a d-dimensional ground state of class s and turn it into a $(d+1)$-dimensional ground state of class $s+1$. So, let there be s pseudo-symmetries $\widetilde{J}_1, \ldots, \widetilde{J}_s$ and a vector bundle $\{\widetilde{A}_k\}_{k \in M_d}$ of class s over d-dimensional momentum space M_d. Our first step is to jack up these data by applying the $(1,1)$ periodicity lemma. Thus, doubling the number of bands, we extend the Clifford algebra of pseudo-symmetries by two generators, I and K, one of which is imaginary (K), and we lift the vector bundle isomorphically to the doubled band space.

Then, introducing a coordinate t for the extra dimension to be added, we make the fibers A_k (of the given vector bundle) t-dependent by applying a one-parameter group of unitary transformations:

$$A_{k,t} := \mathrm{e}^{(t/2)KJ(A_k)} \cdot A_k \,,$$

where $J(A) = \mathrm{i}(\Pi_A - \Pi_{A^c})$ is multiplication by i on A and $-\mathrm{i}$ on the orthogonal complement A^c. By investing the algebraic properties at hand, one verifies that this definition has all the right properties: (i) the Fermi constraint is satisfied; (ii) the pseudo-symmetry relations hold; and (iii) the fibers degenerate at $t = \pm\pi/2$.

Thus we have added one pseudo-symmetry given by the real generator I, and we have increased the space dimension by one. The finished product is a vector bundle $\{A_{k,t}\}_{(k,t) \in M_{d+1}}$ of class $s+1$ over the *suspension* of $M_{d+1} \equiv \widetilde{S}M_d$.

In case you haven't noticed it, let me stress that our Diagonal Map is quite analogous to the Bott map in that it assigns to a point in a symmetric space a minimal geodesic in a related symmetric space. The difference is that our construction has to be compatible with the Fermi constraint, and we have to add one symmetry (not remove one symmetry).

5.4. Diagonal Map versus Bott Map. — The big question now is whether our Diagonal Map of vector bundles yields an isomorphism in homotopy. The answer is: Yes, under certain conditions. I cannot give any details here, but will just flash a picture [drawing] and say this: the Diagonal Map composed with the projection map of a certain fibration yields the Bott Map, after adjusting the real structure. The fibration map essentially doubles the geodesic distance from a certain base point. In this sense one may say that the Diagonal Map is a square root of the Bott Map.

6. Summary

To summarize, this was a story in three parts: first, some current physics of great interest and activity; second some ancient mathematics of great foundational value; and, third, the fascinating story of how the two come together. Last but not least, I need to give major credit to my former PhD student Ricardo Kennedy. The proof that the Diagonal Map is a homotopy equivalence required some serious homotopy theory, and this was done mostly by him.

Here [foto] is Ricardo on the occasion of receiving the Cologne faculty prize for the best PhD thesis in 2016. The PhD advisor was out of town, so he sent the co-advisor instead.

July 24, 2017

MARTIN R. ZIRNBAUER