1 Introduction

This is the continuation of the "Photons, Schmotons" thread, edited by your truly. Part I may be found here:

`http://math.ucr.edu/home/baez/photon/`

—Michael Weiss, editor

2 Michael Weiss: Our story so far

So that's the deal with coherent states for a *single* harmonic oscillator! Maybe a recap would be useful, before whizzing on to Part II.

Our story so far:

Norton T. Noton (no degrees of freedom, scarcely a photon) lives on a line, masquerading as a function $\psi(x)$. When he's feeling especially low, he dresses himself in "basic black": the ground-state function $\psi(x) = (\text{constant})e^{-x^2/2}$. His energy is a mere 1/2 on these occasions, and he refuses to be pinned down, neither as to where he is, nor as to where he's going. A casual observer might think he's going nowhere at all, but the more perceptive sort notice that he's going through a phase— in fact, phase $e^{-it/2}$ in time t, so it takes 4π seconds for him to become his former self again.

Peering at the label on his outfit, we find the cryptic symbol $|0\rangle$.

Give him a good shove, and things start to happen. He acquires momentum b and position c, and bedecks himself in an extravagant corkscrew cape: $\psi(x) = (\text{constant})\exp(ibx-(x-c)^2)$, labelled $e^{-icp+ibq}|0\rangle$. So energetic does he become that before our eyes, he seems to split into two notons— or is it three— or dozens— or hundreds... dissolving, coalescing, impossible to count, a tribute to the art of cinematography...

Our quirky quantum $e^{-icp+ibq}|0\rangle$ has become a coherent superposition: $C_0|0\rangle + C_1|1\rangle + C_2|2\rangle + \ldots$ (where $|n\rangle$ labels a happy clan of n notons). The C_n are proportional to $(z/\sqrt{2})^n/\sqrt{n!}$, where $z = c + ib$. $|C_n|^2$ is proportional to the probability that n notons cavort— or as much cavorting as one can do, with no degrees of freedom! The $|C_n|^2$ form a Poisson distribution with mean value $|z|^2/2$.

The camera pulls back for a wide angle shot. The corkscrew $\psi(x)$ looks more and more like a mere bump on the line, a lump of a particle harmonically sliding back and forth. And in double-exposure overlay, a Poisson distribution of notons fades into the closing credits.

3 John Baez: From the Symplectic to the Complex

OK, class, today we pass from the classical to the quantum, or in other words, from the symplectic to the complex.

A symplectic vector space serves as the "phase space" for a classical system; a complex Hilbert space serves as the "space of states" of a quantum system. Suppose you have a classical field theory and you want to quantize it. Then here's what you can do: first take the phase space of your classical field theory and turn it into a complex Hilbert space using the trick described below. This Hilbert space is the space of "single-particle states" of the quantum field theory you're after. Then you form the "Fock space" on that Hilbert space — a trick which let you describe collections of arbitrarily many particles. Eventually we will do this for the electromagnetic field and get photons. We'll warm up for this using classical systems with finitely many degrees of freedom.

So: how do we turn a symplectic vector space into a complex Hilbert space?

It's not really complicated if you go at it correctly. You start with a real vector space V, real dimension $2n$, with symplectic structure ω. You then introduce a complex structure J on V, i.e. a linear map

$$J : V \to V$$

with $J^2 = -1$. This makes V into a complex vector space, complex dimension n, with the operator J corresponding to multiplication by i.

Then you try to define a complex inner product on V with ω as its imaginary part, in the obvious way:

$$\Im\langle u, v\rangle = \omega(u, v)$$

$$\Re\langle u, v\rangle = \omega(u, Jv)$$

(where $\Im$ means imaginary part and $\Re$ means real part). The second equation is to make sure that $\Re\langle u, v\rangle = \Im\langle u, Jv\rangle$ as it should be.

Now: if we try to define a complex inner product $\langle u, v\rangle$ this way, we may or may not succeed. It depends on J. Some J's are "good" and some are "bad". To be "good", J must have two properties.

First, it needs to preserve the symplectic structure. In other words, we need

$$\omega(Ju, Jv) = \omega(u, v)$$

for all u and v. The reason is that we need our inner product to satisfy

$$\langle Ju, Jv\rangle = \langle u, v\rangle.$$

Second, it must be positive. In other words, we need $\omega(v, Jv)$ to be nonnegative for all v, so that the inner product of any vector with itself is nonnegative, the way it should be.

In general there are a lot of "good" complex structures for a given symplectic structure. Which one is right depends on the physics.

The moral: say we have a classical mechanics problem where the phase space is a symplectic vector space. To quantize, we need to make a somewhat arbitrary choice: a choice of a "good" complex structure. In many cases, the choice is unique if we require it to be invariant under the group of symmetries of the problem, and also require that energies work out to be positive.

4 John Baez: in a more mathematical way in order to confuse you

Wilbert Dijkhof asks:

> I asked Michael what a symplectic structure is. He gave an answer, but I have some more questions.

It seems like you answered all your own questions perfectly well. I'll just summarize: a symplectic structure on a real vector space V is a nondegenerate antisymmetric bilinear form

$$\omega : V \times V \to \mathbf{R}$$

In other words, ω eats pairs of vectors and spits out numbers, and it's bilinear:

$$\begin{aligned} \omega(au + bu', v) &= a\omega(u, v) + b\omega(u', v) \\ \omega(u, av + bv') &= a\omega(u, v) + b\omega(u, v') \end{aligned}$$

antisymmetric:

$$\omega(u, v) = -\omega(v, u)$$

and nondegenerate:

$$\omega(u, v) = 0 \text{ for all } v \text{ implies } u = 0$$

Note that this is very much like a *metric* on V, say

$$g : V \times V \to \mathbf{R}$$

A metric is also required to be bilinear and non degenerate. The only difference is that the metric is symmetric instead of antisymmetric:

$$g(u, v) = g(v, u)$$

and it's nonnegative:

$$g(u, u) \geq 0$$

This makes it easy to describe a symplectic structure in terms of a metric (or vice versa). Namely, if we start with a real vector space V equipped with a metric g, we can define a symplectic structure ω by:

$$\omega(u, v) = g(u, Av)$$

as long as $A : V \to V$ is a 1–1 skew-adjoint operator, that is one for which

$$g(u, Av) = -g(v, Au) \quad \text{(skew-adjointness)}$$

and

$$Av = 0 \text{ implies } v = 0 \quad \text{(1–1-ness)}$$

No surprise here — I just wrote down the conditions we need for ω to work out to be a symplectic structure!

In particular, if V is good old $\mathbf{R}^{2n}$ and g is the good old dot product:

$$g(u, v) = u \cdot v$$

then A can be any 1–1 skew-symmetric matrix. For example, we can take A to be what you suggest:

$$A = \begin{pmatrix} 0 & I \\ -I & 0 \end{pmatrix}$$

where I is the $n \times n$ identity matrix. This gives us a symplectic structure on $\mathbf{R}^{2n}$ which is secretly the same as the imaginary part of the inner product on $\mathbf{C}^n$! Why? Well, let's work it out! This will also help us straighten out the minus signs that invariably screw things up in this business.

We can think of a vector u in $\mathbf{C}^n$ as a pair of vectors u_1, u_2 in $\mathbf{R}^{2n}$ — corresponding to the position and momentum of a particle in n-dimensional space. The inner product to two such vectors in $\mathbf{C}^n$ is then

$$\langle u, v \rangle = u_1 \cdot v_1 + u_2 \cdot v_2 + i(u_1 \cdot v_2 - u_2 \cdot v_1)$$

so the imaginary part of the inner product — our symplectic structure — is

$$\begin{aligned} \omega(u, v) &= u_1 \cdot v_2 - u_2 \cdot v_1 \\ &= (u_1, u_2) \begin{pmatrix} 0 & I \\ -I & 0 \end{pmatrix} \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} \end{aligned}$$

where the last line is matrix multiplication of a row vector, a $2n \times 2n$ matrix, and a column vector. In more civilized notation, we have

$$\omega(u, v) = u \cdot Av$$

where A is the matrix you wrote down.

So in short, you understand completely what's going on, and I'm just restating it in a more mathematical way in order to confuse you.

Now, over on the "supergeometry" thread, we've been talking about how the difference between bosons and fermions amounts to sticking in minus signs whenever you switch two fermions. So one might guess that the difference between inner products and symplectic structures has a lot to do with the difference between bosons and fermions. It's true!

The beautiful thing is that a *complex* inner product has a metric as its real part and a symplectic structure as its imaginary part. In this thread, we're talking about quantizing bosons (e.g. photons), so we're starting with a classical phase space that's a symplectic vector space, and then we're choosing a complex structure — a way to multiply by i. This lets us define a metric, and when we put all the ingredients together we get a complex Hilbert space — the "single-particle space" of our quantized theory.

But if we were talking about quantizing fermions, things would be slightly changed. We would start with a classical phase space that's a real vector space with a metric,

and then we'd choose a complex structure. This would let us define a symplectic structure, and then when we put all the ingredient together we'd again get a complex Hilbert space. This would again be the "single-particle space" of our quantized theory.

So: complex Hilbert spaces have a bosonic aspect to them (the symplectic structure) and a fermionic aspect to them (the metric), tied together inextricably by the complex structure! It's cool how the math and the physics fit together so well.

One more thing. When you hear the word "metric", you might think of general relativity. And you might make a wild guess: "Well, if a vector space with a metric wants to be the classical phase space of a *fermion*, could there be some fermionic aspect to gravity?"

Well, it's not quite so simple. But if you dig into it really deep and straighten out all the problems with this idea, you *do* get something out of it — the spin network/spin foam approach to quantum gravity. (A good place to start is actually Penrose and Rindler's *Spinors and Spacetime*.)

5 Michael Weiss: a very nice picture, but a little too abstract.

OK, let's look at a special case, 1-D quantum mechanics.

The classical 1-dimensional phase space is a 2-dimensional vector space, so it looks like a plane. We give it coordinates (q, p).

The symplectic form is defined on pairs of vectors in this phase space. It's defined by:

$$\omega((q_1, p_1)(q_2, p_2)) = \begin{vmatrix} q_1 & p_1 \\ q_2 & p_2 \end{vmatrix} = q_1 p_2 - p_1 q_2$$

so it's just the area of the parallelogram spanned by the vectors (q_1, p_1) and (q_2, p_2).

The Hamiltonian $(q^2 + p^2)/2$ defines a vector field on phase space via Hamilton's equations, which I won't write out for fear of making a sign error. But the vector field is easy to picture: little arrows swirling around the origin in the clockwise direction (with our sign conventions). So that the Hamiltonian flow just rotates the plane clockwise at a uniform rate. Points in phase space trace out circles, all with the same period.

Notice that the Hamiltonian flow is area-preserving, as we expect from Liouville's theorem.

The case of the harmonic oscillator is quite special (to paraphrase the Church Lady). In general the configuration space is a manifold, and the phase space is the cotangent bundle to this manifold. But here configuration space is a vector space (namely the real line), so the cotangent bundle is a two-dimensional vector space. Flat as Texas (ever see that great poster, "Ski Lubbock!"?) Generalizing to n-dimensions is no sweat; dropping flatness is much more significant (as I understand things).

OK, now we turn our 2-dimensional vector space into the complex plane. We decree that

$$i(q, p) = (-p, q)$$

and so if we identify $(q, 0)$ with the real number q, we can write $(q, p) = q + ip$. (You know, I think it was Hamilton who first *defined* complex numbers as pairs of real numbers.)

This is what John means by giving phase space a complex structure. The key is defining what it means to multiply an element of phase space by i.

If we started with n oscillators, our configuration space would be $\mathbf{R}^n$. So phase space would be $\mathbf{R}^{2n}$. That would morph into $\mathbf{C}^n$, i.e., n-dimensional complex Hilbert space.

Returning to 2-dimensional Flatland, or rather 1-complex-dimensional Hilbert space $\mathbf{C}$... the inner product of two vectors (a) and (b) is a^*b, i.e., the conjugate of a times b. (Or is it ab^*? Another convention to screw up.)

Or spelling out the real and imaginary parts:

$$\langle (q_1 + ip_1), (q_2 + ip_2) \rangle = (q_1 q_2 + p_1 p_2) + i(q_1 p_2 - p_1 q_2)$$

The imaginary part of this inner product is $\omega((q_1, p_1)(q_2, p_2))$. So I guess we did want a^*b. As John wrote it:

$$\Im \langle u, v \rangle = \omega(u, v)$$

$$\Re \langle u, v \rangle = \omega(u, Jv)$$

Note also that the real part is the ordinary Euclidean dot product in two real dimensions.

Our ω-preserving Hamiltonian flow has now changed into the inner-product-preserving map:

$$q + ip \mapsto (q + ip)e^{-it}$$

and the jargon-meisters rush forward plastering "symplectomorphism" over "ω-preserving", and "unitary" over "inner-product-preserving".

This is a very nice picture, I admit, but a little too abstract. Why are we doing this? What's the punch-line?

I've chewed this to a fine mush. Here's the deal: in the classical picture, we have a few interesting observables, all functions from phase space to the reals:

q,	aka position
p,	aka momentum
$H = \frac{\text{radius}^2}{2}$,	aka energy

Turning to the quantum picture, as a Hilbert space, $\mathbf{C}$ is about as bare-bones as it gets. It sure isn't the Hilbert space we *usually* think about when we hear the words, "harmonic oscillator"! *That* Hilbert space is L^2, and has some nice fun-guy operators living in it, like our old pals the creation and annihilation operators (over there by the video set, playing Doom.)

6 John Baez: a little too wet

Michael Weiss writes:

> OK, now we turn our 2-dimensional vector space into the complex plane.

Right.

> This is a very nice picture, I admit, but a little too abstract.

Too abstract, eh? Reminds me of the guy who said the ocean was very nice — but a little too wet.

True, we need to work quite a bit before these abstractions come to life for us. Perhaps the main thing you have to wrap your mind around is the analogy between quantum mechanics and quantum field theory. In both of these subjects, you start with a "little" Hilbert space and then form the Fock space on that, which is "big". The "little" Hilbert space is the same as the phase space of the classical system,

while the "big" Hilbert space is the Hilbert space for the corresponding quantum system.

For example: the classical phase space of a single particle on the line is $\mathbf{C}$. The Hilbert space for the corresponding quantum system is the Fock space on $\mathbf{C}$, which has a basis of states $|n\rangle$. This space can also be thought of as $L^2(\mathbf{R})$. $\mathbf{R}$ is the classical *configuration* space of a single particle on the line. If we "complexify" the configuration space $\mathbf{R}$ we get the classical phase space $\mathbf{C}$.

Or: the classical phase space of a vibrating string is some Hilbert space H. A point in phase space describes both the position and momentum of the string. Thus elements of H are certain pairs of real-valued functions on the interval $[0, 1]$. The Hilbert space for the corresponding quantum system is the Fock space on H, which has a basis of states

$$|n_1, \ldots, n_k\rangle.$$

This space can also be thought of as $L^2(X)$, where X is the classical *configuration* space of a single particle on the line. An element of X is a single real-valued function from $[0, 1]$. If we "complexify" the configuration space X we get the classical phase space H.

Or: start with the classical phase space H of a real scalar field on n-dimensional space, and do everything in analogy to the case of a string. Elements of H are pairs of real-valued functions on $\mathbf{R}^n$.

Or: start with the classical phase space for electromagnetism on n-dimensional space, and do the same game. Elements of H are pairs consisting of a "vector potential" and "electric field".

Here is the twist that is probably confusing you. Since in any of these examples the classical phase space is a Hilbert space, we can *also* think of it as the Hilbert space of a *quantum* system. In the quantum field theory examples, states in the Fock space describes collections of particles, while the classical phase space is the Hilbert space for a *single particle*. But in the quantum mechanics examples, the classical phase space is distressingly small!

As you note:

> Turning to the quantum picture, as a Hilbert space, $\mathbf{C}$ is about as bare-bones as it gets. It sure isn't the Hilbert space we *usually* think about when we hear the words, "harmonic oscillator"!

What does the Hilbert space $\mathbf{C}$ mean here? Well, it's the Hilbert space of a "single quantum with no internal degrees of freedom" — whatever *that* means. This boring old quantum has but a single state (apart from the all-important phase).

Then when we form the Fock space over $\mathbf{C}$, we get a space with a basis of states like $|n\rangle$. The state $|n\rangle$ describes a collection of n "quanta with no internal degrees of freedom".

So: in this example the "Hilbert space for a single particle" is utterly dull, and only when we form the Fock space do we get some fun annihilation and creation operators.

In the more fancy examples, there is already some interest in the Hilbert space for a single particle.

But the thing to keep in mind is that the "Hilbert space for a single particle" is also the "classical phase space".

7 Michael Weiss: A finer mush

Lessee, last time I said:

> I've chewed this to a fine mush.

But the mush wasn't quite fine enough, here's a finer mush.

Our old friend, the classical harmonic oscillator, we can picture nicely as a dot racing on a circle. More precisely, the classical *state* of the oscillator is given by the position of this dot.

Now let's go over to quantum-land. $\mathbf{R}^2$ becomes $\mathbf{C}$. Doesn't *seem* like such a big deal. Aha, but $\mathbf{C}$ is *not* the set of quantum *states*; it's the Hilbert space of *state-vectors*!

That's the key. To get the quantum states, we have to "mod out" by the complex numbers. You know how it goes: state-vectors v and cv are regarded as equivalent for any non-zero complex number c. (Here v must also be non-zero. So technically we're forming equivalence classes of the Hilbert space minus the zero vector.)

So what does this "modding out" process do to our poor old classical states? Well, for the 1-D harmonic oscillator, it completely kills off any individuality.

First off, forget the energy observable! Doesn't matter how far out you are— doesn't matter which circle in phase space is your race-track— you're all equivalent. In other words, we can normalize the state-vector v, replacing it with $v/\|v\|$.

That makes a weird kind of sense. $\mathbf{C}$ is the Hilbert space for a *single* quantum "particle"— a single "noton". A single noton always has energy $\hbar/2$. To get more energy, you need more notons. We saw how that worked with coherent states.

What about the position observable? Normalizing our state-vector puts it on the circle of radius 1. But in classical-land, it still gets to race around! The state-vector v becomes $e^{-it}v$. Different classical states, but the *same* quantum state.

And this also makes quantum sense. If you're in an eigenstate of the Hamiltonian— if you have definite energy— you don't change! That ain't so in classical-land. In quantum-land, life changes only for *superpositions* of energy eigenstates. We've seen how that works with coherent states too.

So I guess it's a three step plan:

1. Start with a classical phase space, full of sound and fury and sword-play.
2. Turn it into a Hilbert space. This is the home for a *single* quantum "particle". But what a drab world our quantum particle inhabits! States that classically had different energies have become the same. Phase-space trajectories that classically "took you somewhere" now just spin you in place.
3. "Second quantize": create the Fock space, the home of *many* quantum particles. Form coherent superpositions. And the richness is recovered.

Step (2) hinges on the choice of a complex structure. It tells us what $\|v\|$ is, and what $e^{-it}v$ is.

To sum it up in a diagram:

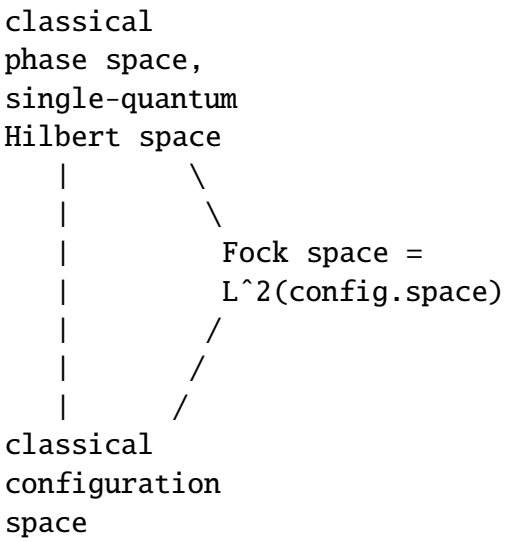

8 John Baez: The choice of complex structure

Let me write (q, p) for a point in the phase space for a particle on the line. We can also think of this as the point $q + ip$ on the complex plane.

Then the symplectic structure ω is imaginary part of the usual inner product of complex numbers:

$$\begin{aligned} \omega((q,p),(q',p')) &= \Im\langle q + ip, q' + ip' \rangle \\ &= p'q - q'p \end{aligned}$$

while the complex structure J corresponds to multiplication by i:

$$J(q,p) = (-p,q) \qquad \text{since } i(q+ip) = -p + iq$$

Beware of minus signs: we can get different formulas with different signs if we write $(p, q) = p + iq$ instead of $(q, p) = q + ip$, or work with an inner product that is linear in the first slot rather than in the second slot. Conventions vary.

However, we shouldn't worry *too* much about the conventions! We should keep our eyes on the main point: the classical mechanics of a linear system only requires that the phase space be a symplectic real vector space, while its quantum mechanics (e.g. constructing the Fock space) requires that it be a complex Hilbert space. To get from the former to the latter we need to pick a good complex structure, where "good" means that $\omega(Ju, Jv) = \omega(u, v)$ for all u and v, and $\omega(v, Jv)$ is nonnegative for all v. The above example is the simplest example of how this works. However, there are different choices for J, even in this simple example.

The need for this choice causes various subtleties in quantum field theory. For example, when doing quantum field theory on curved spacetime there is a canonical way to choose a symplectic structure on the space of classical solutions of the wave equation, but no canonical way to choose a complex structure. This makes the notion of "Fock space" — and thus the notions of "vacuum" and "particle" — somewhat arbitrary. Hence Hawking radiation.

Alternatively, say we are studying particles coupled to a external gauge field. There is again an obvious best symplectic structure, but different external fields define different complex structures. So the notion of "particle" depends in a nontrivial way on the external field. Worse, as we trace out a nontrivial loop in the space of external fields, our notion of "particle" can fail to come back to where it was! There can be a funny sort of phase ambiguity, called an "anomaly". This

is usually regarded as a bad thing, so one throws out gauge field theories where the particles have anomalies. People predicted the existence of the top quark because if one quark in a given generation were simply missing, there would be an anomaly.

9 Michael Weiss: A particle with personality!

Class is now in session.

So, what's this on the board?

> Find the symplectic structure for the wave equation in 1 dimensional space. Find the unique Poincaré-invariant complex structure on the space of solutions of the (real!) wave equation.

I had a lot of trouble with this problem, and all because I violated Wheeler's Moral Principle: "Never calculate anything until you know the answer!" So instead of giving the answer outright, I think I will first mull out loud about it for a bit.

OK, the wave equation: $u_{xx} = u_{tt}$, where the subscripts stand for partial derivatives: $u_{xx} = \frac{\partial^2 u}{\partial x^2}$, $u_{tt} = \frac{\partial^2 u}{\partial t^2}$. (Here we set the wave propagation speed equal to 1, of course.) We want to look at the space of all real-valued solutions to this.

For starters we have the classical configuration space. Easy enough to picture: a violin string. Pair u up with $\partial u/\partial t$ and we have the classical phase space of the string: where it is, and where it's going. Slap on a complex structure and we've got the Hilbert space for a single particle. A particle with personality! — with the tonal richness of a Stradivarius, or at least a one-string fiddle: let's call it a *foton*.

Violin strings have been a hot topic from 1747, when D'Alembert first derived the wave equation, to 1900, when Fejér finally proved rigorously that the Fourier series of a continuous periodic function is Cesàro summable to that function. (And beyond: important new theorems were still appearing as late as 1966. My analytical friend Carl has written up a fascinating historical sketch, from which I learned all this.)

We want to look at *harmonically oscillating* modes of the string. Perhaps I should say "listen to" instead of "look at"! These morph into eigenstates of the foton: states where the foton just basically twiddles its phase, *not changing*. States like this, for example:

$$u(x,t) = \cos(kx - \omega t) \quad \text{(A travelling wave)}$$

Plugging into the wave equation $u_{xx} = u_{tt}$ we get $k^2 = \omega^2$, or $|\omega| = |k|$. (You'll have to keep straight when ω means angular frequency and when it means symplectic structure.)

Umm. *Real* violin strings— even *real ideal* violin strings— have fixed ends. In technical jargon, boundary conditions: say, $u(0,t) = u(1,t) = 0$ for all t. But that would rule out our simple travelling wave. Of course we could use periodic boundary conditions.... I think I'll postpone that topic to another time. For now, our violin string is *infinitely long*.

Lessee, what's our screenplay here? Configuration space → phase space → Hilbert space, right? Last time we watched the saga of a lowly classical harmonic oscillator which, through hard work and sacrifice, achieved the status of a lowly "noton", a particle with no degrees of freedom. The one-note soundtrack went like this:

configuration space	$q(t) = \cos t$
phase space	$(q(t), p(t)) = (\cos t, -\sin t)$
Hilbert space	$q(t) + ip(t) = \exp(-it)$

This time we start with a travelling wave $u(x,t) = \cos(kx - \omega t)$, and I guess the score is:

configuration space	$u(x,t) = \cos(kx - \omega t)$
phase space	$(u(x,t), \dot{u}(x,t)) = (\cos(kx - \omega t), \omega \sin(kx - \omega t))$
"Hilbert space"	$\exp(i(kx - \omega t))$

—except for the embarassing fact that $\exp(i(kx - \omega t))$ has "infinite norm", and isn't really a citizen-in-good-standing of a true Hilbert space. Let's just ignore that little problem, shall we?

I've used $\dot{u}$, the time-derivative of u, for the "momentum". Seems like a good bet.

The answer has just the right form: $\exp(i(kx - \omega t)) = e^{-i\omega t} \exp(ikx)$, so it is a rotating vector in "Hilbert space", and the corresponding quantum state-vector *doesn't change*. Not only that, but it's wonderfully cinematic! This complex wave function looks like an infinite string that wraps around the line corkscrew fashion. It turns in place, and waves appear to travel along it, like the stripes on a barber-pole.

So are we ready to answer the home-work question? For notons, we had $(q, p) \leftrightarrow q + ip$. Now we want:

$$(\cos(kx - \omega t), \omega \sin(kx - \omega t)) \leftrightarrow e^{i(kx - \omega t)} = \cos(kx - \omega t) + i \sin(kx - \omega t)$$

Writing $q = \cos(kx - \omega t)$, $p = \omega \sin(kx - \omega t)$ to reduce the clutter a bit:

$$(q, p) \leftrightarrow q + \frac{i}{\omega} p$$

This gives us the rule for translating between phase space and "Hilbert space". Multiplying both sides by i, on the left we should get the long-sought complex structure on our phase space:

$$J(q, p) \leftrightarrow -\frac{1}{\omega} p + iq$$

Using the translation rule on the right hand side:

$$(-\frac{p}{\omega}, \omega q) \leftrightarrow -\frac{1}{\omega} p + iq$$

or finally:

$$J(q, p) = (-\frac{p}{\omega}, \omega q)$$

We've defined the complex structure! The rest is just mopping up.

Umm. Quite a *lot* of mopping up. Poincaré invariance, symplectic structure, the energy, the *meaning* of ω for any old configuration of the string, and of course those boundary conditions.... This is beginning to look more like the aftermath of a Steven Segal action-flic, and not the bio-pic of that celebrated classical violinist, Frances P. Foton! But I think that's enough for now.

10 John Baez: I don't know much Yiddish, but...

Michael Weiss wrote:

> A particle with personality! — with the tonal richness of a Stradivarius, or at least a one-string fiddle: let's call it a *foton*.

Gee, you're making up particles like a true physicist: first schmotons, then the one-note notons, and now fotons! (Actually, I was the one who corrected the barbaric "shmoton" to read "schmoton". I don't know much Yiddish, but terms of derogation seem to need that c — when a schlemiel starts schmoozing, or a schmuck starts his schtick, that extra c invariably comes in. Besides, if we get rid of the c, why keep the h — might as well set them *both* equal to 1, leaving us with "smotons"?)

configuration space	$q(t) = \cos t$
phase space	$(q(t), p(t)) = (\cos t, -\sin t)$
Hilbert space	$q(t) + ip(t) = \exp(-it)$

Right. Here, of course, time evolution just amounts to rotation an angle t clockwise. (Clockwise? Well, that's because e^{-it} goes around clockwise. This has to do with our convention of writing points in phase space as (q, p) rather than (p, q). But after all, isn't it quite fitting for time evolution to go clockwise?) So in particular, the complex structure, which is multiplication by i, corresponds to time evolution *backwards* a quarter period — the "backwards" bit being an unfortunate spinoff of our conventions. But let's just keep that in mind...

> This time we start with a travelling wave $u(x, t) = \cos(kx - \omega t)$, and I guess the score is:

configuration space	$u(x, t) = \cos(kx - \omega t)$
phase space	$(u(x, t), \dot{u}(x, t)) = (\cos(kx - \omega t), \omega \sin(kx - \omega t))$
"Hilbert space"	$\exp(i(kx - \omega t))$

Looks good. What happens if we do time evolution backwards a quarter period in this case? Do we get the complex structure you worked out?

> Quite a *lot* of mopping up... This is beginning to look more like the aftermath of a Steven Segal action-flic

Actually it's beginning to look like an Irving Segal paper on quantum field theory! But let's see, can you show the complex structure is invariant under space and time translations, for a start?

11 Michael Weiss: Fools fear to tread...

Last time I offered a "definition" for a complex structure on our "violin string" space, which I'll call V. Recall the definition of V: pairs of real-valued functions $(u, \dot{u})$, representing initial conditions for solutions to the classical wave equation $u_{tt} = u_{xx}$. (As before, I use subscripts to indicate partial derivatives.) We could also think of the elements of V as *being* these solutions: given a pair $(u(x), \dot{u}(x))$, we find the unique function $u(x, t)$ for which:

$$\begin{aligned} u_{tt} &= u_{xx} \\ u(x,0) &= u(x) \text{ for all } x \\ u_t(x,0) &= \dot{u}(x) \text{ for all } x \end{aligned}$$

—at least we can do that if $u(x)$ and $\dot{u}(x)$ are nice enough. I'm being deliberately vague about boundary conditions and other such persnickety details; after all, gotta give der Herr Professor JB *something* to write about! (Last time I used infinitely long violin strings; that is, I pushed the boundaries off to infinity. Some musicians prefer periodic boundary conditions, or in the lingo, they "put those fotons in a box".)

So here was my "definition" of a complex structure J on V: if

$$\begin{aligned} q &= \cos(kx - \omega t) \\ p &= \omega \sin(kx - \omega t) \end{aligned}$$

then q is a solution to the wave equation if $k^2 = \omega^2$. Also $p = q_t$, so (q, p) at $t = 0$ is one of our $(u, \dot{u})$ initial conditions. Define:

$$J(q,p) = (-\frac{p}{\omega}, \omega q)$$

Well, I haven't really defined J on V, have I! I've defined J for *certain elements* of V. Or at least if you'll let me sweep the boundary conditions under the edge of the rug! (Where *else* would you sweep boundary conditions?)

But my set of (q, p) pairs ought to get us pretty far, especially if we throw in the complementary set:

$$\begin{aligned} q &= \sin(kx - \omega t) \\ p &= -\omega \cos(kx - \omega t) \end{aligned}$$

also with $k^2 = \omega^2$. If $t = 0$, then we have the makings of a Fourier series, or a Fourier integral (again depending on those pesky boundary conditions), so we ought to be able to extend my definition of J "by linearity" to *all* of V.

Consider it done. Fools fear to tread where die Herren Professoren rush right in, so I'll leave the funky analytical details for someone who really knows what they're doing.

Still, the overall tune is really quite pleasant. Thanks to this complex structure, we can think of our first $q(x,t)$ as the real part of the complex function $e^{i(kx-\omega t)}$. Also $p(x,t)$ corresponds to the imaginary part, but with that factor of ω thrown in.

I will voice a worry or three I've been having about this business. The usual Fourier integral formula:

$$\psi(x,t) = (\text{constant}) \int \hat{\psi}(\omega,k) e^{i(kx-\omega t)} dk$$

doesn't look very Poincaré invariant. (Here I innocently set $\omega = |k|$, so I don't have to integrate over *both* k and ω. I have a sneaking suspicion we will hear more about that.) Poking around in the textbooks a bit (Haag's *Local Quantum Physics*) suggests replacing this with:

$$\psi(x,t) = (\text{constant}) \int \hat{\psi}(\omega,k) e^{i(kx-\omega t)} dk/|k|$$

What does this do to the Fourier inversion formula? Or to our formula for the complex structure?

(Incidentally, Poincaré invariance might be a good reason *not* to impose periodic boundary conditions. Kind of hard to have a Poincaré invariant box— even a music box!)

And what about the inner product? We were promised a Hilbert space, once we found a complex structure!

I was kind of hoping that the "vectors" $e^{i(kx-\omega t)}$ would form a Dirac-style "orthonormal basis" for this Hilbert space. This doesn't seem to quite work out...

Oh well. More unfinished business. John Baez added two more items to the list. I'll tackle just the first one in the remainder of this post.

> What happens if we do time evolution backwards a quarter period in this case? Do we get the complex structure you worked out?

No. But that's OK. I'd just worked out *why* this is fine, and was trying to illustrate it, when I ran across a neat applet that shows the difference perfectly. Go to Michael Fowler's on-line physics notes and take a gander at:

`http://landau1.phys.Virginia.EDU/classes/109N/more_stuff/GroupVelocity.html`

Fowler doesn't *say* he's illustrating the difference between the complex structure and time evolution; as the URL suggests, the applet depicts group velocity and phase velocity. A multi-purpose applet: set the group velocity to zero, and the applet illustrates the complex structure; set it to one to one, and the applet illustrates time evolution.

How's that? Well, let's take a simple wave— just not *too* simple. Let's take $u(x, 0)$ to be the sum of two pure harmonic waves, say:

$$u(x, 0) = \cos(k_1 x) + \cos(k_2 x)$$

If k_1 and k_2 are very close, then $u(x, t)$ looks like wiggly scribble confined between two slowly varying cosine waves.

If we want to time-evolve u, we just let t increase a bit, taking us from $u(x, 0)$ to $u(x, t)$:

$$u(x, t) = \cos(k_1 x - \omega_1 t) + \cos(k_2 x - \omega_2 t)$$

If instead we want to apply J, we think of $u(x, 0)$ as the real part of

$$\psi(x) = e^{ik_1 x} + e^{ik_2 x}$$

and multiply ψ by i. A little more generally, we could multiply ψ by *any* phase factor, say e^{-ic}. This is a kind of "complex structure" evolution. The parameter c takes the place of t, and we get:

$$\psi(x, c) = e^{i(k_1 x - c)} + e^{i(k_2 x - c)}$$

So with time-evolution, the phase of each term changes by a *different* amount, because t is multiplied by different values of ω. With complex-structure evolution, each phase changes by the *same* amount.

The group velocity is defined as

$$d\omega/dk$$

or for our two-note wave,

$$\Delta\omega/\Delta k = (\omega_1 - \omega_2)/(k_1 - k_2)$$

But $\omega = |k|$ for our fotons; when k_1 and k_2 are positive, the group velocity is 1. So plug that value into Michael Fowler's applet, and you've got your picture for time evolution.

Complex-structure evolution looks like time-evolution, except with ω set to the constant value 1 for all k. Which means that the group velocity will be zero.

12 John Baez: the covariant phase space formalism

Michael Weiss wrote:

> Last time I offered a "definition" for a complex structure on our "violin string" space, which I'll call V.

Yes indeed, and it was very cleverly arrived at: you simply thought of a violin string as a collection of independent harmonic oscillators, and used the standard complex structure on the phase space of each one of these harmonic oscillators. That's what harmonic analysis is all about: when in doubt, diagonalize!

> Recall the definition of V: pairs of real-valued functions $(u, \dot{u})$, representing initial conditions for solutions to the classical wave equation $u_{tt} = u_{xx}$. (As before, I use subscripts to indicate partial derivatives.) We could also think of the elements of V as *being* these solutions.

Right. That goes by the name of the "covariant phase space formalism": we think of a point in phase space as *being* a solution of whatever equation we're messing with. The advantage of this over the older approach — where we think of a point in phase space as being *initial data* for a solution — is that it doesn't require us to make an arbitrary choice of time to call $t = 0$. This choice becomes a bigger deal in special relativity, where different people have different ideas of t, and even more so in general relativity. The moral: to avoid an arbitrary split of spacetime into space and time, work with the covariant space formalism! Then, whatever groups of symmetries happen to be around will act in an *obvious* way on your phase space.

There is however a price to pay: it's sometimes harder to write down explicit formulas in the covariant phase formalism! Sometimes we want to get our grubby little hands on some concrete p's and q's, or in the present situation, u's and $\dot{u}$'s. For that, we need to choose a time to call "now" — the famous $t = 0$ or "spacelike slice" that canonically minded physicists always begin the day with.

But there's no real conflict: we can work using the covariant phase space formalism whenever that's convenient, and work with initial data whenever *that's* convenient.

> I'm being deliberately vague about boundary conditions and other such persnickety details; after all, gotta give der Herr Professor JB *something* to write about!

So you get to do all the juicy stuff, leaving the dry details to me? Harrumph! (he said Germanically).

But just so people realize these details are not really so scary, let me sketch out how one goes about filling them in.

The initial data u and $\dot{u}$ don't need to be particularly nice to determine a unique solution of the wave equation. If you want nice, let them be arbitrary smooth functions; if you want nasty, let them be arbitrary "distributions", like the nineteenth derivative of the Dirac delta function. It all depends how good you are at *interpreting* the wave equation in a sufficiently general way.

For smooth initial data, we interpret the equation

$$u_{tt} = u_{xx}$$

just the way our calculus teachers told us to — you know, epsilons and deltas and all that jazz. And it's easy to solve: we just *guess* the solution, using the fact that every solution is a sum of a left-moving wave and a right-moving wave, each trucking along at unit speed:

$$u(x,t) = f(t-x) + g(t+x).$$

If we work with distributions as initial data we need to interpret the derivatives in the distributional sense — but don't worry, they teach you want that means when you learn what the heck "distributions" are! And again, the wave equation is easy to solve: you just guess

$$u(x,t) = f(t-x) + g(t+x)$$

and figure out what f and g have gotta be to match the initial data u and $\dot{u}$.

Anyway, there's no big problem solving the darn thing given some initial data. As for why there's a *unique* solution, well, just assume there are two, subtract them, and get a contradiction... I leave the details to you. (Hint: conservation of energy.)

So we can start out leaving things pretty flexible, and decide *exactly what kind of solutions* we're talking about later, when we construct our complex Hilbert space of solutions. This Hilbert space will, as usual, have a norm, and we will then restrict attention to "finite-norm" solutions — i.e., solutions that lie in the Hilbert space! One standard way to proceed is this: first we take the space of compactly supported smooth initial data, and then complete this with respect to the norm (which we will eventually have a formula for) to get the precise Hilbert space we're after.

But anyway, this is only important if you want to dot every i and cross every t. For now let's sort of be relaxed about the functional-analytic aspects and focus on the physics.

> Last time I used infinitely long violin strings; that is, I pushed the boundaries off to infinity. Some musicians prefer periodic boundary conditions, or in the lingo, they "put those fotons in a box".

Putting your fotons in a box doesn't refer to *periodic* boundary conditions in particular; it just refers to working with waves that stay inside a box-shaped region of space, and if course if one does so one needs *some* boundary conditions to say what happens to the waves when they hit the wall of the box. Periodic boundary conditions are mathematically the easiest to deal with, but they're a little odd, since they say the waves "wrap around" and immediately appear on the other side. A mirrored box of light or violin string would be better dealt with using some other boundary conditions, like "Dirichlet" boundary conditions, since the waves reflect off the wall, not "wrap around".

Anyway, if we aren't completely exhausted by the time we answer your original question about photons, we might consider photons in a box (or maybe watered-down "fotons") and calculate the magnitude of the Casimir effect, which is due to the vacuum energy and depends on the choice of boundary conditions.

But I digress... I think I'll continue in another post.

13 John Baez: Multiplication by i on a Thursday

In his last post, Michael fleshed out the details of his proposed complex structure on the phase space of solutions of the wave equation in 2 dimensions. He call this phase space V and thinks of points in it as pairs $(u, \dot{u})$ consisting of our solution u at time zero and its first time derivative $\dot{u}$. These are supposed to be analogous to pairs (q, p) consisting of the position and momentum of a point particle. And the

idea is to use this analogy to figure out everything about the wave equation, starting from what we already know about the harmonic oscillator. A plucked string is just a bunch of uncoupled harmonic oscillators corresponding to the different vibrational modes of the string! Each one has its own little 2d phase space with its own complex structure. We glom them all together to get the complex structure on V.

By the way, don't forget why we're doing this! Our goal is to make V into a complex Hilbert space, so we can start doing quantum field theory: unit vectors in V will describe states of a single quantum particle — a "foton" — when we quantize the wave equation. However, V started out life as a classical phase space, and is thus a real vector space equipped with a symplectic structure. The symplectic structure will be the imaginary part of the inner product once we succeed in making V into a complex Hilbert space, but to get the real part of the inner product, and to make V into a complex vector space in the first place, we need a way to multiply vectors in V by i — a complex structure!

Remember how the complex structure worked for the harmonic oscillator. A point in the phase space for the harmonic oscillator is just a pair of real numbers:

$$(q, p)$$

We can think of a pair of real numbers as a single complex number, and then the complex structure — multiplication by i — gives the point (q, p) a quarter turn counterclockwise. To emphasize that we are really *introducing* a notion of multiplication by i, instead of using a pre-given one, we write the complex structure as J instead of i. So we have:

$$J(q, p) = (-p, q).$$

Now, all this may seem obvious, if perhaps rather odd, but it's really important not to merely our heads and say "yup, that's multiplication by i, all right" — we have to understand what's *right* about this definition of J. First of all, it satisfies

$$J^2 = -1$$

which is just what it means for a linear operator to be a complex structure. But there are *lots* of different operators we could have picked with *that* property. Secondly, it preserves the symplectic structure: we have

$$\omega(Jv, Jv') = \omega(v, v')$$

where $v = (q, p)$ and $v' = (q', p')$ are points in our phase space and the symplectic structure ω is given by

$$\omega(v, v') = pq' - qp'$$

But there are *lots* of complex structures that also have *this* property. Why? See Hint 1 below if you can't figure it out. Thirdly, it commutes with time evolution. This is important, since we don't want multiplication by i on Thursday to be different from multiplication by i on Friday! Remember that time evolution for the harmonic oscillator is simply rotation: more precisely, to evolve in time by an amount t, we rotate the phase space t radians clockwise. (Our standard harmonic oscillator has period 2π.) So our complex structure also has to commute with rotations! And it turns out that there are only *two* complex structures that commute with rotations: the one we are calling J, and also $-J$. (If you don't see why, try Hint 2.)

So the question is just: why is J better than $-J$? Well, this depends on certain conventions, but having chosen these conventions already, the answer is: because the Hamiltonian in the *quantum* theory of the harmonic oscillator should be positive! Time evolution in quantum mechanics is usually given by

$$e^{-itH}$$

where H is a *positive* operator. In the case at hand, if we take i to be our operator J, we see that $H = 1$, which is positive. If we took i to be $-J$, we would get $H = -1$. That would be negative, which we decree to be bad.

This stuff is important. The reason Dirac thought negative- frequency solutions of the Dirac equation corresponded to negative energy states of the electron, and thus had to make up a fancy story about positrons being "holes" in the "sea of negative energy states", was that he got the wrong complex structure!

And if Michael isn't careful, we're gonna get the wrong complex structure here, too! Of course, I will put all the blame on him if this happens — I hereby absolve myself of all responsibility! But really, he should be able to do it right. The great thing about minus signs is that, unlike most things in life, you always have at least a 50% chance of getting them right!

So let's see how Michael is doing it. A vibrational mode of our infinite string is characterized by a frequency ω and a wavenumber k, which describe how fast our string wiggles in the t and the x directions, respectively. Corresponding to each frequency ω and wavenumber k with $\omega^2 = k^2$, we have *two* linearly independent solutions of the wave equation:

$$u(x,t) = \cos(kx - \omega t)$$

and

$$u(x,t) = \sin(kx - \omega t)$$

So we get a little 2-dimensional phase space, just as we should — in the end, everything should look just like our good old harmonic oscillator, but with some ω's and k's thrown in to spice things up. Remember, that's the basic principle of quantum field theory: "Approximate the universe by a bunch of harmonic oscillators; everything else is too complicated".

And so, what Michael appears to be doing, is to define a complex structure on this little 2-dimensional phase space. That seems good — we can worry later about glomming all these little phase spaces together to get V. But the big question is: does Michael's complex structure J commute with time evolution? If so, it has a 50% chance of being right! It's right if in addition, the Hamiltonian works out being positive.

It's possible that Michael has the right J when the frequency ω is positive, and wrong when it's negative. That's the mistake Dirac made.

Also, last time I asked:

> What happens if we do time evolution backwards a quarter period in this case? Do we get the complex structure you worked out?

and Michael answered:

> No. But that's OK.

to which I can only reply: Hmm. Are we really on the same wavelength here, or are we talking at cross-purposes? I've shown above that for the harmonic oscillator, there are only two complex structures that commute with time evolution: time evolution a quarter period *forwards* in time, and time evolution a quarter period *backwards*. So your formula for the complex structure on your little 2-dimensional phase space, to be reasonable, should be one or the other of those! Of course what counts as "a quarter period" depends on ω, so it depends on *which* little phase space we're at.

Hint 1: For the harmonic oscillator, a linear transformation of the phase space preserves the symplectic structure if and only if it's *area-preserving*. The usual choice of J — a quarter-turn counterclockwise — clearly has this property. But we can conjugate J by any area-preserving map T and get another complex structure $J' = TJT^{-1}$ that also has this property!

Hint 2: The only linear transformations of the plane that commute with all rotations are combinations of rotations and dilations. There aren't many such transformations with square -1. Just two, in fact!

14 Michael Weiss: A mere difference in viewpoint

John Baez wonders, in that tell-tale tone:

> to which I can only reply: Hmm. Are we really on the same wavelength here, or are we talking at cross-purposes?

(Long-time readers of JB will know that this "Hmmm" can only mean, "I doubt the sanity of thee or me, and sooner thee than me.")

I had, you see, answered his question:

> What happens if we do time evolution backwards a quarter period in this case? Do we get the complex structure you worked out?

with a nonchalant:

> No. But that's OK.

But then he added casually:

> Of course what counts as "a quarter period" depends on ω, so it depends on *which* little phase space we're at.

Whew! A mere difference in viewpoint— no need to check myself into Bellevue just yet. Sure, if you restrict attention to a single little phase space, then time-evolution and what I called "complex structure evolution" are the same.

But I'm trying to build up an album of mental pictures and video clips— sort of a "That's QFT, Folks!"— and some of these should illustrate the modding out

process. (Recall that we get the set of quantum states from the Hilbert space by "modding out the complex numbers": v and cv give the same quantum state for any non-zero complex number c.)

I can picture a violin string OK. But then, that's not *quite* the same as the quantum states of a single foton.

If v is a vector in one of our little phase spaces— say, v is the pair of functions

$$(\cos(kx - \omega t), \omega \sin(kx - \omega t))$$

then v time-evolves to $e^{-it\omega}v$, which describes the *same* foton state.

So I picked the next simplest case after this: the superposition of vectors from *two* little phase spaces, $v_1 + v_2$. Since the ω's are different, the corresponding foton state *really changes* over time.

On the other hand, if we multiply $v_1 + v_2$ by a non-zero complex number, the foton state *doesn't* change— by definition!

Then I stumbled across Fowler's applet that pictures this— icing on the cake. (Check it out if you haven't already:
`http://landau1.phys.Virginia.EDU/classes/109N/more_stuff/GroupVelocity.html`.)

Roughly speaking, it seems like the foton state corresponds to the *envelope* of the function $v_1 + v_2$. The superposition looks like this at $t = 0$ (before we "mod out"):

$$u(x, t) = \cos(k_1 x) + \cos(k_2 x)$$

which (thanks to the wonders of trigonometry) is the same as:

$$u(x, t) = 2 \cos\left(\frac{k_1 + k_2}{2} x\right) \cos\left(\frac{k_1 - k_2}{2} x\right)$$

Loosely speaking, the "foton state" corresponds to the second factor, $\cos(\frac{k_1 - k_2}{2} x)$. The foton has a definite "beat frequency", and even a "beat phase", but the phase of the fast "inner wiggles" is indeterminate. (If the quoted phrases don't make sense, just look at the applet.)

I'd *also* like to get a better mental picture of a real Gaussian wave-packet, travelling right, let's say. What does *that* look like after you apply J to it?

15 John Baez:

> Whew! A mere difference in viewpoint— no need to check myself into Bellevue just yet.

Okay, good. I *figured* we were just talking at cross-purposes.

> Sure, if you restrict attention to a single little phase space, then time-evolution and what I called "complex structure evolution" are the same.

Right. More precisely, just to wrap it up in a neat package: for any eigenstate of the Hamiltonian, time evolution by a quarter period backwards in time is multiplication by i!

(And of course this is true for *any* quantum system, not just the sort we are considering.)

Okay, so where are we? We're trying to quantize the wave equation in 2-dimensional Minkowski spacetime:

$$u_{tt} = u_{xx}$$

where $u(x, t)$ is a real-valued function of one space coordinate x and one time coordinate t. When we succeed in doing this, we'll have the quantum field theory of "massless neutral spin-0 particles": massless because the wave equation describes waves moving at the speed of light, spin-0 because u is just a scalar field, and neutral because it's real-valued. (For a charged particle we'd use a complex-valued u.)

Or, if the language of particle physics is too scary, we can alternatively say we are giving the quantum description of an idealized violin with a single very, very long string.

But right now we're trying to do the very first step: to make the space of solutions into a complex Hilbert space. So far we've been working mostly "one vibrational mode at a time". Corresponding to each vibrational mode there's a little 2-dimensional phase space, having as a basis the two solutions

$$\cos(kx - \omega t)$$

and

$$\sin(kx - \omega t)$$

for a fixed wavenumber k and frequency ω satisfying $k^2 = \omega^2$. And we've shown how to equip this little phase space with a symplectic structure and a complex structure, thereby making it into a 1-dimensional complex Hilbert space. (2 real dimensions equal 1 complex dimension.)

But now, how do we glom all these little 1-dimensional complex Hilbert spaces together, to get the complex Hilbert space of *all* solutions of the wave equation? Well, there is a standard technique for glomming together a lot of Hilbert spaces: it's called taking a "direct sum", or in the present case, where we are glomming together a continuous family of Hilbert space, a "direct integral". Everyone should know this technique. It makes the next step a snap.

But, alas, not everyone knows about direct integrals, so I will instead use the Fourier transform, which everyone *does* know. [1]

And in fact Michael has already sketched the basic game plan...

First, we notice that any reasonable function $u(t, x)$ can be written as

$$u(t, x) = \frac{1}{2\pi} \int \hat{u}(\omega, k) e^{i(kx - \omega t)} d\omega \, dk$$

for some function $\hat{u}$ called the Fourier transform of u. [2]

Now, it's easy to see that if u satisfies the wave equation then $\hat{u}$ must be zero except where $k^2 = \omega^2$, and conversely. [3] Note that the equation $k^2 = \omega^2$ describes an "X"-shaped locus in Fourier transform space, which is called a lightcone. We think of it as being made of two parts:

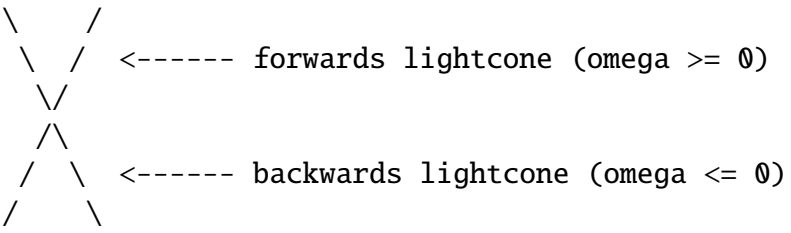

It's also easy to see that if u is real, $\hat{u}$ satisfies

$$\hat{u}(\omega, k)^* = \hat{u}(-\omega, -k),$$

where the * denotes complex conjugation. And conversely.

Thus to have an arbitrary real solution u of the wave equation is the same as to have an arbitrary complex function $\hat{u}$ on the forwards lightcone!

Why the forwards lightcone? Because the equation above says that the value of $\hat{u}$ on the forwards lightcone determines its value on backwards lightcone.

So the obvious *guess* for how to make the space of real solutions of the wave equation into a complex Hilbert space is to use L^2 of the forwards lightcone — square-integrable functions on the forwards lightcone. Of course, we have to pick some measure on the forwards lightcone for this to make sense — you can't do an integral without a measure. Luckily there is a unique measure on the lightcone that's invariant under Lorentz transformations, namely $dk/|k|$. [4] If we use that, then by general abstract nonsense the Lorentz transformations will automatically act as unitary operators on L^2 of the forwards lightcone. And translations in spacetime *obviously* act as unitary operators (I claim), so we'll be happy: the whole Poincaré group will act as unitary operators on our Hilbert space.

This obvious guess works out to be right, and to be consistent with the stuff we have already done "one vibrational mode at a time". But maybe we should work it out in detail. ("We" meaning Michael, of course.) For example, here's one thing we could do: show that the obvious complex structure on L^2 of the forwards lightcone — the usual notion of multiplication by i — agrees with the complex structure Michael already worked out "one vibrational mode at a time".

> I was kind of hoping that the "vectors" $e^{i(kx-\omega t)}$ would form a Dirac-style "orthonormal basis" for this Hilbert space. This doesn't seem to quite work out...

Well, the most important reason it doesn't quite work out is that the functions $e^{i(kx-\omega t)}$ aren't *real*, and our Hilbert space only contains *real* solutions of the wave equation. (There are other, lesser reasons too.)

You could at this point ask why we didn't study *complex* solutions of the wave equation. And the reason is that this raises temptations that are very difficult to resist. Namely, when you are studying *complex* solutions to some linear equation, there's a blatantly obvious complex structure staring you in the face: the usual multiplication by i. It's very tempting to use this, but it's usually *wrong* — the Hamiltonian won't work out to be positive. Dirac fell for this temptation and got very confused. He then started waving his arms frantically and muttering about "holes" in the "sea". By means of such handwaving he eventually did something that amounted to picking the *right* complex structure. But unfortunately few physicists even to this day realize that one can avoid the problem in the first place if one is careful to pick the right complex structure — not the obvious one.

Since we're studying *real* solutions of the wave equation, it's more clear that we have to do something clever to give our Hilbert space of solutions a complex structure. Note how we're doing it: we work in the Fourier transform picture with $\hat{u}$ instead of u, and then the trick is to multiply $\hat{u}$ by i the forwards lightcone and multiply it by $-i$ on the backwards lightcone! We *have* to multiply it by $-i$ on the backwards lightcone since

$$\hat{u}(\omega, k)* = \hat{u}(-\omega, -k)$$

for the Fourier transform of a real function!

Okay, I'm tired out for now... I'm too tired to comment on Michael's stuff about group vs. phase velocity, except to say that it's cool. Also, I have no idea what our complex structure does to a solution of the wave equation that looks like a Gaussian, except to note that Gaussians have an incredible tendency to become other Gaussians when you do things to them.

Notes:

[1] Actually the Fourier transform is just a special case of doing a direct integral: we are taking the space $L^2(\mathbf{R})$ and expressing it as an integral over of the 1-dimensional Hilbert spaces spanned by the functions e^{-ikx}. Of course, the functions e^{-ikx} do not themselves lie in $L^2(\mathbf{R})$, but the direct integral is designed to deal with precisely this sort of thing. I'm not sure what's the best place to learn about direct integrals, but I seem to recall a lot of discussion of them in Dixmier's *Von Neumann Algebras* and Takesaki's *Theory of Operator Algebras I*.

[2] I am starting to write $u(t, x)$ instead of $u(x, t)$ since in particle physics that's how they do it. Similarly I'll write $\hat{u}(\omega, k)$ with ω coming before k.

[3] So in fact $\hat{u}$ is a distribution, not a function. To be honest, we should say: the Fourier transform of any tempered distribution is again a tempered distribution.

See Reed and Simon's *Fourier Analysis and Self-Adjointness* or some such tome. We're secretly gonna write $\hat{u}$ as some function times some measure living on the lightcone, and then we're gonna call this function $\hat{u}$ to cover our tracks.

[4] Well, okay, it's unique up to a scalar factor. If we already happen to know the right symplectic structure on the space of solutions of the wave equation, we can use this to determine the right scalar factor, using the fact that the imaginary part of the inner product in our L^2 space had better be the symplectic structure. Right now I'm hoping the scalar factor is 1.