WHY A TIMPANI IS SO HARD TO TUNE

ARAV BHATTACHARYA

Introduction

In high school, I was a percussionist. And, as a percussionist, I was stood at the back of the orchestra, very rarely having to tune my instruments. When I did, it was often just to adjust the timbre of my snare drum with a drum key. Now, compare that with my peers at the front of the orchestra playing violin. Every day, the violinists would make everyone's ears bleed with their tuning before finding their A440.

I remember not even thinking about tuning when I started. When I told my band director that the marimba was out of tune, I learned that we didn't tune it because the wooden bars would literally have to be shaven with some kind of power tool. More embarrassingly, I only realized that timpani heads needed to be tuned when I literally saw my band teacher tightening the lug nuts that secured the drum head and asked what she was doing.

Now, my band teacher loved having us students take responsibility for things. Us percussionists were usually tasked with keeping our equipment in shape, and this included tuning our concert snare and replacing its drum heads. So why didn't she tell me, or any one of the other percussionists, to tune the timpani?

Initial thought: maybe the drums didn't have to be tuned all that often, and so she just decided to take care of it herself. A quick Google search disproved this: a self-proclaimed percussion teacher on Reddit claimed that they adjusted the tuning of their timpanis daily.

So there's some deeper reason at play. As any reasonable person would do, I next turned to mathematics to see if I could find this deeper reason. Surprisingly, this proved productive: I found a mathematical reason for why timpanis were harder to tune accurately than violins. For the rest of this essay, I'll take you down the same rabbit hole I went down to find the answer.

First, we'll visit the dynamics of violins and timpanis, and derive the wave equation whose solutions dictate the movement of the strings and heads of the respective instruments when they are played. Next, we'll develop an 'infinite-dimensional theory of linear algebra' to give us a basis on which we can analyze the solutions of the wave equation. Using our newfound theory, we'll compare the pitches produced by bowing a violin to those produced by hitting a timpani and see how they differ.

Prerequisites. Some familiarity with the definitions of linear algebra and real analysis is necessary to read this essay. The essay builds up any techniques for deriving and solving partial differential equations from first principles, and technical details from the theories of measure and distribution are not discussed in much depth.

Date: October 10, 2025.

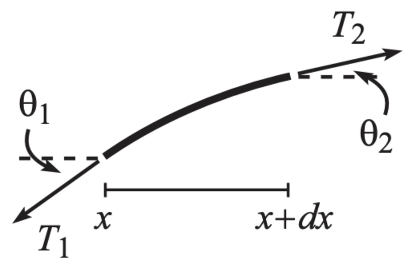

Figure 1. Free-body diagram of an oscillating string (image credit [Mor])

1. The Dynamics of Music

We can start by thinking about the dynamics of idealised violins and timpani[1]. In both cases, these instruments produce sound via the waves produced by a certain medium (a string or drum head, respectively) when it is struck. Hence, we will derive a 'wave equation' that constrains the motion of the waves each respective instrument produces.

Wave Equation, One Dimension

First, we analyze the case of the violin, where the waves produced by its string are one-dimensional. Let $u(x,t)$ be the function denoting the height of the string at location x and time t. By looking closely at this system, we want to find a set of equations specifying the physical constraints that any wave must satisfy.

Let's make clear our assumptions. Let our violin string be of length L with mass density ρ. We know that the ends of the string are fixed, so we say that $u(0,t) = u(L,t) = 0$ (the *Dirichlet boundary condition*). We assume throughout the string, the tension $T \gg F_g$, where F_g is the gravitational force on the string. We will lastly impose no additional constraints on how the string is struck, so that our initial conditions are that the string initially starts with some heights $u(x,0) = u_0(x)$ and is struck with velocity $u_t(x,0) = v_0(x)$.

Now, we know that we are looking for a function $u : [0,L] \times [0,\infty) \to \mathbb{R}$ determining the height of our string. Suppose we had such a function. Picking (x,t) in its domain, we can draw a local free-body diagram of the string near x at time t. This diagram will look something like Figure 1.

Let's consider the net forces acting upon this small segment of the string at time t. First, let's look at the forces in the horizontal direction. We can assume that there is no horizontal displacement of this segment of the string, so that the net force here is $T_2 \cos\theta_2 - T_1 \cos\theta_1 = 0$. If the angle θ is small on both ends of the string segment, in fact we'd have that each of the cosine terms in this net force equation is roughly 1. Hence as the Taylor series for cosine around 0 has error that varies with the square of the angle, we can make the approximation that $T_2 \approx T_1$ (as θ^2 is negligible with respect to θ). We'll call this shared value T, and henceforth take this to be the constant uniform tension at each point on the string.

Now, in the vertical direction, we'd have a net force of $F = T(\sin(\theta_2) - \sin(\theta_1))$. Now, note that the angle θ at a given point is just the value of the slope u_x at that point. Since

[1] In particular, we do not consider sources of energy loss such as air resistance.

we assumed that θ is small, we get that in fact:

$$F \approx T(\theta_2 - \theta_1) = T\big(u_x(x+dx) - u_x(x)\big)$$
$$\rho\, dx\, u_{tt} = T\, dx\, \frac{u_x(x+dx) - u_x(x)}{dx}.$$

The second equation on the previous page is due to Newton's second law, $F = ma$, where in this case $m = \rho\, dx$ is the mass and $a = u_{tt}$ is the acceleration. Now, the rightmost term on the right hand side of this equation approaches u_{xx} as $dx \to 0$ (this can be formalized via the mean value theorem, as u_x is continuously differentiable by assumption) and so after some simplification we get that in the limit

$$u_{tt} = \frac{T}{\rho}\, u_{xx}.$$

We call this final equation the *one-dimensional wave equation*, the equation governing the motion of a one-dimensional wave on a string.

Wave Equation, Two Dimensions

The derivation for the wave equation in the case of a timpani membrane is remarkably similar to the case of the violin string.

Say our timpani has a circular membrane with radius R. We assume that the boundary of the membrane is tightly attached to the body of the drum, giving us Dirichlet boundary conditions $u = 0$. We further assume that the membrane has a constant uniform surface tension $S \gg {F_g}^2$. We denote the membrane's mass density ρ and initial position $u(x, y, 0) = u_0(x, y)$. and we suppose that it is struck with velocity given by $u_t(x, y, 0) = v_0(x, y)$. The main difference this time is that our function has a two-dimensional spatial domain, i.e. we want $u : \overline{B}(0, R) \times [0, \infty)$ where $\overline{B}(0, R) \subset \mathbb{R}^2$.

Let's consider now a small spatial patch near a fixed (x, y), i.e. $[x, x+\Delta x] \times [y, y+\Delta y]$ for small $\Delta > 0$. In particular, we can look at cross-sections of this patch along the y direction.

Now, the force exerted on the ends of the edge of this patch has value $S\, \Delta y$ because the whole patch along the y direction is pulling on this end. As the length of the patch along this direction is Δy, and the unit surface tension is S, we get that the forces are in fact $S\, \Delta y$. Knowing this, we can now draw a cross-sectional free-body diagram as in the 1-D case (see Figure 2).

We can once again use analogous reasoning to the one-dimensional case to get that the force acting on this patch along the y direction is $S\, \Delta x\, \Delta y\, u_{xx}$. By a similar analysis, we find that along the x direction the force on the patch is $S\, \Delta x\, \Delta y\, u_{yy}$. Thus the net vertical force on the patch is the sum of these, and since the mass of the patch is $\rho\, \Delta x\, \Delta y$ and its acceleration is u_{tt} we get that

$$\rho\, \Delta x\, \Delta y\, u_{tt} = S\, \Delta x\, \Delta y\, u_{xx} + S\, \Delta x\, \Delta y\, u_{yy}.$$

We simplify to get the *two-dimensional wave equation*

$$u_{tt} = \frac{S}{\rho}(u_{xx} + u_{yy}).$$

[2]This assumption follows from similar reasoning to the 1-D case above.

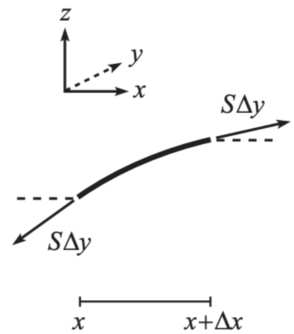

Figure 2. Cross-sectional free-body diagram of an oscillating membrane (image credit [Mor])

Solving the Wave Equation, One Dimension

Let us first disregard the initial conditions and look for solutions that satisfy the boundary conditions. Our solutions to the relaxed problem will later help us find the particular solution to our original problem with given initial conditions.

First, in one-dimension (this solution is adapted from [Oli]). We will set $c = \sqrt{\frac{T}{\rho}}$ to simplify notation. We want to find $u : [0, L] \times [0, \infty) \to \mathbb{R}$ such that $u_{tt} = c^2\, u_{xx}$ and $u(0, t) = u(L, t) = 0$, and we'll soon see why c is called the *wave speed*. We can conjecture that u is a product of a function in x with a function in t, i.e. $u(x, t) = F(x)G(t)$. This conjecture and the associated solution strategy are both called *separation of variables*, and solutions found via separation of variables are called *separable solutions*.

Having separated variables, we via our wave equation see that $F(x)G''(t) = c^2F''(x)G(t)$ and so

$$\frac{F''(x)}{F(x)} = \frac{G''(t)}{c^2G(t)}.$$

But now, the left-hand side of the equation is dependent only on x while the right-hand side is only dependent on t. As the variables x and t are independent, it follows that $\frac{F''(x)}{F(x)}$ and $\frac{G''(t)}{c^2G(t)}$ must both be equal to a constant, say $-\lambda^2 \in \mathbb{R}$ [3].

Rearranging, we see that we've got a second-order ODE for $F(x)$, i.e. F must satisfy $F''(x)+\lambda^2F(x) = 0$. As we know that $u(0, t) = u(L, t) = 0$, we have two possibilities. Either $G(t) = 0$ and our solution is trivially $u \equiv 0$, or we have boundary conditions $F(0) = F(L) = 0$. We examine the latter.

We know that our ODE has solution of form $F(x) = A\sin(\lambda x) + B\cos(\lambda x)$, and the boundary conditions imply that $B = 0$ and $A\sin(\lambda L) = 0$. Assuming nontriviality, $A \neq 0$ so that $A\sin(\lambda L) = 0$ forces $\lambda = \frac{n\pi}{L}$ for some positive integer n. So, we're left with $F(x) = A\sin(\frac{n\pi x}{L})$.

We can similarly solve for $G(t)$ which must satisfy $G''(t) + c^2\lambda^2G(t) = 0$ to get that $G(t) = C\sin(\frac{n\pi ct}{L}) + D\cos(\frac{n\pi ct}{L})$. We would get 'boundary conditions' for this ODE if we imposed initial conditions on the wave equation, which would allow us to solve for specific values of C and D. However, the boundary conditions in this case would depend on x, so

[3] This odd choice of variable will make our lives easier down the road.

that we wouldn't be able to solve for them in the ODE setup. We'll discuss a fix to this problem later in this essay.

For now, we'll just leave these coefficients and say that we have found a set of solutions

$$u_n(x,t) = \sin\left(\frac{n\pi x}{L}\right)\left(a_n \sin\left(\frac{n\pi ct}{L}\right) + b_n \cos\left(\frac{n\pi ct}{L}\right)\right)$$

for nonnegative integer n, $a_n = AC$, and $b_n = AD$ as above. Notice how the value of c scales the period of the wave. This is why we call c the wave speed.

Solving the Wave Equation, Two Dimensions

Now, we turn our attention to the case of two dimensions. The solution strategy is similar here, once again conjecturing that the solution to the wave equation is a product of one-variable functions, i.e., separation of variables.

We first separate t from the spatial variables, i.e. we set $u(x,y,t) = w(x,y)G(t)$. For notational convenience, we define the Laplacian operator $\nabla^2 = \partial_{xx} + \partial_{yy}$. Now, substituting into the wave equation yields

$$w(x,y)G''(t) = c^2\nabla^2 w(x,y)G(t),$$

where $c = \frac{S}{\rho}$ is analogous to before. By rearranging so the t terms are all on the left, we once again get that $\frac{G''}{c^2G}$ is constant, say $-\lambda^2$. We thus get that $G(t) = C\cos(c\lambda t) + D\sin(c\lambda t)$ as before, for some constants C and D.

Plugging this back in to the original wave equation, we are left with the *spatial Helmholtz problem* where

$$\nabla^2 w + \lambda^2 w = 0.$$

We will once again solve this via separation of variables, but first we'll want to make the boundary conditions more elegant. Specifically, if we first change to polar coordinates, the boundary condition simply becomes $w(R,\theta) = 0$, and we even get a bonus condition of 2π-periodicity in θ.

Without further ado, let us convert our two-dimensional wave equation to polar coordinates, i.e. let's change variables $(x,y) = (r\cos\theta, r\sin\theta)$. Omitting a significant amount of algebra, the Helmholtz equation becomes

$$w_{rr} + \frac{1}{r}w_r + \frac{1}{r^2}w_{\theta\theta} + \lambda^2 w = 0.$$

It's now time to separate variables! We set $w(r,\theta) = F(r)G(\theta)$. This yields

$$F''(r)G(\theta) + \frac{1}{r}F'(r)G(\theta) + \frac{1}{r^2}F(r)G''(\theta) + \lambda^2 F(r)G(\theta) = 0.$$

Multiplying both sides by $\frac{r^2}{FG}$ yields

$$r^2\frac{F''(r)}{F(r)} + r\frac{F'(r)}{F(r)} + \frac{G''(\theta)}{G(\theta)} + \lambda^2 r^2 = 0.$$

We can move the angular part of this differential equation to the right, and therefore get that $\frac{G''}{G}$ is constant (say n^2). Therefore, we get that $G(\theta) = A\cos(n\theta) + B\sin(n\theta)$. To maintain the 2π-periodicity of $G(\theta)$, we will also need that $n \in \mathbb{Z}$.

Plugging G back in and multiplying both sides by F, we get that

$$r^2F''(r) + rF'(r) + \left(\lambda^2 r^2 - n^2\right)F(r) = 0.$$

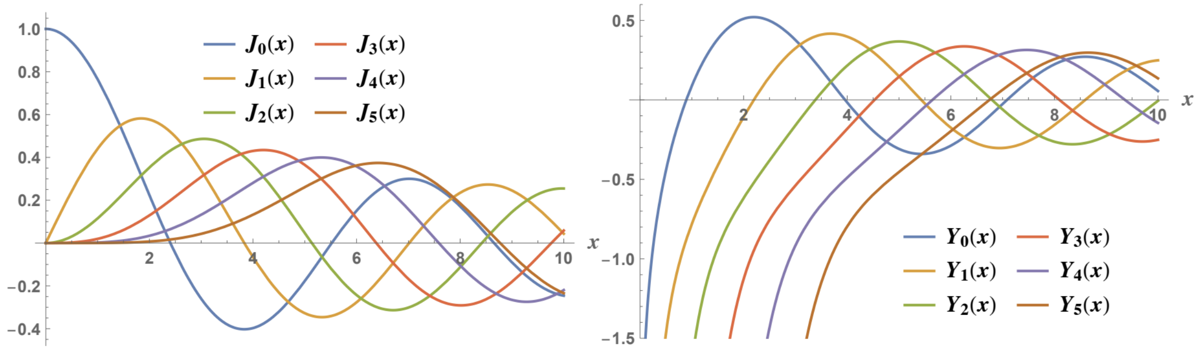

(a) (b)

Figure 3. (a) Bessel functions of the first kind J_n. (b) Bessel functions of the second kind Y_n (image credit [Hew])

We perform one last substitution $z = \lambda r$ and since $F_r = \lambda F_z$ and $F_{rr} = \lambda^2 F_{zz}$ we get that $F(z)$ must satisfy the *Bessel equation*

$$z^2 F''(z) + zF'(z) + \left(z^2 - n^2\right) F(z) = 0.$$

The solutions to the Bessel equation are called Bessel functions. For a given value of n in the equation above, these functions are denoted J_n (Bessel functions of the first kind) and Y_n (Bessel functions of the second kind). More information about these functions, including their exact formulas, can be found in [Hew]. For now, we simply plot the graphs of the first few of these functions in Figure 3.

Thus we get that $F(r) = CJ_n(\lambda r) + DY_n(\lambda r)$. As Y_n has a singularity at $r = 0$, but J_n and our function u do not, we must have $D = 0$. Furthermore, since $F(R) = 0$, we must have that λR is a zero of the Bessel function J_n, say $j_{n,m}$. Then $\lambda = \frac{j_{n,m}}{R}$ and so we have found solutions $u(r, \theta, t)$ to the two-dimensional wave equation via separation of variables. Specifically, we've found u to be as follows:

$$\begin{aligned} u(r,\theta,t) = &\left[a_{n,m}^{(c)} \cos\left(\frac{cj_{n,m}\, t}{R}\right) + b_{n,m}^{(c)} \sin\left(\frac{cj_{n,m}\, t}{R}\right)\right] J_n\left(\frac{j_{n,m}\, r}{R}\right) \cos(n\theta) \\ &+ \left[a_{n,m}^{(s)} \cos\left(\frac{cj_{n,m}\, t}{R}\right) + b_{n,m}^{(s)} \sin\left(\frac{cj_{n,m}\, t}{R}\right)\right] J_n\left(\frac{j_{n,m}\, r}{R}\right) \sin(n\theta). \end{aligned}$$

How Do We Solve the Equation with Initial Conditions?

Having solved relaxed versions of the wave equations that constrain how our instruments should move when producing sound, we might ask how these relaxed solutions can help us solve the wave equation given the initial conditions. Before we move any further, let's unify notation and write the wave equation as $u_{tt} = c^2 \nabla^2 u$ regardless of dimension. In particular, we'll define the Laplacian operator $\nabla^2 = \partial_{xx}$ in the one-dimensional case. Additionally, in both cases, let Ω be the relevant domain for our function u.

Now, here comes the fun part. We can notice that in fact, the set of the solutions to each of these PDEs form a vector space! Why so? Well, $C^2(\Omega)$, the set of twice continuously differentiable functions on Ω, is a vector space. The set of possible solutions $\mathcal{S}$ is a subset of $C^2(\Omega)$. So we just need to perform a quick subspace test to verify that the functions solving the wave equation form a vector subspace. Evidently, $u \equiv 0 \in \mathcal{S}$ by the homogeneity of the wave equation. Furthermore, given two solutions $u_1, u_2 \in \mathcal{S}$, $au_1 + bu_2 \in \mathcal{S}$ for any $a, b \in \mathbb{R}$ by its linearity. Hence the set of solutions to the wave equation is indeed a vector space.

What's even nicer is that, with suitable coefficients, the functions we've found above are orthonormal under the inner product

$$\langle f, g \rangle = \int_\Omega fg,$$

which the reader may want to verify.

We might hope to find some clever way to prove that these orthonormal functions are not just a linearly independent set, but a basis. Establishing this fact would prove that we have found a basis for this solution space.

Given that these functions are also eigenfunctions of the Laplacian operator ∇^2, maybe we could apply the spectral theorem? If we could find some clever way to apply the spectral theorem, then we'd then be able to infer that the eigenvalues of these functions are all the eigenvalues of ∇^2 on the domain Ω. These eigenvalues would correspond to resonant frequencies produced by the instrument, in a very precise manner[4]. We can then express any physically possible vibration of each instrument as being a superposition of the eigenfunctions. In fact, we'd be able to go a bit further, and say that given some initial conditions, there is a unique superposition of eigenfunctions of the wave equation determining the vibration of the instrument.

Unfortunately, the fact that all the solutions we found above are linearly independent functions means that our vector space of solutions is infinite-dimensional. So, as much as the above situation seems like an obvious application of the spectral theorem to this vector space, our finite-dimensional linear algebra is not a strong enough theory for us to pursue our intuition in this infinite-dimensional setting. Hence, we must turn to an infinite-dimensional version of linear algebra, better known as functional analysis.

In the next chapter, we will develop this theory of functional analysis, which will ultimately enable us to finish our mathematical description of violins and timpanis.

2. Linear Algebra in Infinite Dimensions

Before we actually jump into functional analysis, let's take a moment to discuss why our finite-dimensional linear algebra is insufficient here.

In finite dimensions, our first instinct when proving statements is often to use arguments that involve induction on the basis elements. Unfortunately, in infinite dimensions, we cannot exhaustively go through the now infinitely large basis. Worse still, it can be shown that the classical definition of a basis, wherein each element of a vector space must be a finite linear combination of scaled basis elements, means that every basis of an infinite-dimensional space must be uncountably large (see [Car05]). Thus there's no hope of induction at all!

To rectify this problem, we'll introduce a relaxed definition of a basis, allowing for infinite sums of basis elements. But this introduces even more challenges! In finite dimensions, we also take for granted being able to represent elements of our vector space via coordinate representations. But now, these coordinate representations will have infinitely many coordinates, which naturally leads to questions about convergence of the underlying series of basis elements. Taking hints from analysis, we will also rectify this problem.

[4] Specifically, the frequencies will turn out to be the square roots of these eigenvalues.

Then, we'll work our way towards the spectral theorem, which will work just as in finite dimensions to allow us to decompose the solution space $\mathcal{S}$ into the direct sum of eigenspaces, and just as in finite dimensions give us an orthonormal basis that we can express solutions of the wave equation under. We'll conclude by showing that the orthonormal bases correspond to the solutions we found in the previous section.

Schauder Bases

Let's start by recalling some definitions from finite dimensions, and see how we can amend them to fit our infinite-dimensional setting.

First, let's rectify our problem with bases. Recall the definition of a basis from finite dimensions:

Definition 2.1. Let V be a vector space over the field $\mathbb{F}$ and suppose $\mathcal{B} \subseteq V$. We say that $\mathcal{B}$ is a *Hamel basis* of V if every element $v \in V$ can be expressed as

$$v = \sum_{i \in I} a_i e_i$$

for some finite indexing set I, $a_i \in \mathbb{F}$, and $e_i \in \mathcal{B}$.

The reason that we are calling bases by our original definition 'Hamel bases' is in order to differentiate them from our new notion of a Schauder basis.

Definition 2.2. Let V be a vector space over the field $\mathbb{F}$ and suppose $\mathcal{B} \subseteq V$ is a countable set, and let I be a set indexing the elements e_i of $\mathcal{B}$. We say that $\mathcal{B}$ is a *Schauder basis* if every element $v \in V$ can be expressed as

$$v = \sum_{i \in I} a_i e_i$$

for some $a_i \in \mathbb{F}$.

Note that our main change is that we can now express vectors as countable linear combinations of elements of the Schauder basis. We restrict to countable linear combinations so that the sum is well-defined without going into the uncountable realm and dealing with measure theory.

The other change is that we also don't allow the whole basis to be uncountable. As you might conjecture, this means that not every vector space even has a Schauder basis[5]. This is another key difference from Hamel bases, because every vector space has a Hamel basis.

Now having updated our definition of a basis, we will find a more restricted definition of vector spaces that gives us nice convergence properties.

Hilbert Spaces

The sharp-eyed reader will notice that for convergence of infinite series to make any sense, we must have some additional topological structure on a vector space. We fix this issue by exclusively working on inner product spaces for the remainder of this essay. We'll denote the inner product $\langle \cdot, \cdot \rangle$ and the norm induced by the inner product $\| \cdot \| = \sqrt{\langle \cdot, \cdot \rangle}$.

We now turn our attention to the issue of convergence. What matters to us is convergence of sequences in terms of the norm induced by the inner product. To enforce that the

[5]See [Kre91] for more details, including why Hilbert spaces are the only inner product spaces with Schauder bases.

convergence behaves as we would intuitively expect, we can ask that our space is complete in the induced norm, which in fact is our next definition:

Definition 2.3. A *Hilbert space* is an inner product space that is complete in the norm induced by the inner product.

Hilbert spaces inherit some very nice geometrical properties from inner product spaces, and we'll now turn our attention to a few that will be useful in later proofs.

Lemma 2.4 (Cauchy-Schwarz and triangle inequalities)**.** *The Cauchy-Schwarz inequality states that in an inner product space, the norm induced by the inner product satisfies*

$$|\langle x, y\rangle| \leq \|x\|\|y\|,$$

where equality holds exactly when $\{x, y\}$ is a linearly dependent set.

The inner product norm also satisfies

$$\|x + y\| \leq \|x\| + \|y\|,$$

where equality holds when there is some nonnegative real c such that $x = cy$ or $y = cx$.

Proof. We begin with the Cauchy-Schwarz inequality, and use it to prove the triangle inequality. Notice that the Schwarz inequality holds trivially for $y = 0$, since $\langle x, 0\rangle = 0$. Thus we take $y \neq 0$. For every scalar α we have

$$\begin{aligned} 0 \leq \|x - \alpha y\|^2 &= \langle x - \alpha y, x - \alpha y\rangle \\ &= \langle x, x\rangle - \overline{\alpha}\langle x, y\rangle - \alpha(\langle y, x\rangle - \overline{\alpha}\langle y, y\rangle). \end{aligned}$$

Taking $\overline{\alpha} = \langle y, x\rangle / \langle y, y\rangle$, we simplify our inequality to

$$\begin{aligned} 0 &\leq \langle x, x\rangle - \frac{\langle y, x\rangle}{\langle y, y\rangle}\langle x, y\rangle - \alpha(\langle y, x\rangle - \frac{\langle y, x\rangle}{\langle y, y\rangle}\langle y, y\rangle) \\ &= \langle x, x\rangle - \frac{\langle y, x\rangle}{\langle y, y\rangle}\langle x, y\rangle - \alpha(\langle y, x\rangle - \langle y, x\rangle) \\ &= \langle x, x\rangle - \frac{\langle y, x\rangle}{\langle y, y\rangle}\rangle = \|x\|^2 - \frac{\langle y, x\rangle\,\langle x, y\rangle}{\|y\|^2} \\ &= \|x\|^2 - \frac{|\langle x, y\rangle|^2}{\|y\|^2}. \end{aligned}$$

Rearranging, we get that $|\langle x, y\rangle|^2 \leq \|x\|^2\|y\|^2$, and taking square roots on both sides gives us the Cauchy-Schwarz inequality.

Now, the equality cases require either $y = 0$ or $\|x - \alpha y\|^2 = 0$, in which case $x = \alpha y$. In either case, $\{x, y\}$ is linearly dependent.

Now, let's move on to the triangle inequality. By the Cauchy-Schwarz inequality,

$$|\langle x, y\rangle| = |\langle y, x\rangle| \leq \|x\|\|y\|.$$

Thus we have that

$$\begin{aligned}\|x+y\|^2 &= \langle x+y, x+y\rangle \\ &= \|x\|^2 + \langle x,y\rangle + \langle y,x\rangle + \|y\|^2 \\ &\le \|x\|^2 + 2|\langle x,y\rangle| + \|y\|^2 \\ &\le \|x\|^2 + 2\|x\|\|y\| + \|y\|^2 \\ &= (\|x\|+\|y\|)^2.\end{aligned}$$

Taking square roots on both sides, we get the triangle inequality.

We have equality here exactly when

$$\langle x,y\rangle + \langle y,x\rangle = 2\|x\|\|y\|.$$

Note that the left-hand side of this equation is $2\Re\langle x,y\rangle$, where $\Re$ denotes the real part. By the Schwarz inequality, we must then have

$$\Re\langle x,y\rangle = \|x\|\|y\| \ge |\langle x,y\rangle|.$$

Since the real part of any complex number cannot exceed its absolute value, we must have that both sides of the above Cauchy-Schwarz inequality are equal, whereby x, y are linearly dependent; furthermore, the imaginary part of $\langle x,y\rangle$ must be zero, so $\langle x,y\rangle$ is a positive real number.

Without loss of generality, say we have $x = cy$ for nonzero y (the only case this does not cover is $y = 0$, which follows by switching the variables). We will show that c is a real nonnegative number. Since we have that

$$0 \le \Re\langle x,y\rangle = \Re\langle y,y\rangle = |\langle cy,y\rangle| = c\|y\|^2,$$

we can divide both sides by $\|y\|^2$ to get that $c \ge 0$. ∎

Another quick insight into the geometry of inner product spaces is the parallelogram equality:

Lemma 2.5 (parallelogram equality)**.** *In an inner product space, we have that*

$$\|x+y\|^2 + \|x-y\|^2 = 2(\|x\|^2 + \|y\|^2).$$

Proof. We have that

$$\begin{aligned}\|x+y\|^2 + \|x-y\|^2 &= \langle x+y, x+y\rangle + \langle x-y, x-y\rangle \\ &= \langle x,x\rangle + \langle x,y\rangle + \langle y,x\rangle + \langle y,y\rangle + \langle x,x\rangle - \langle x,y\rangle - \langle y,x\rangle + \langle y,y\rangle \\ &= 2\|x\|^2 + 2\|y\|^2 = 2(\|x\|^2 + \|y\|^2).\end{aligned}$$

∎

The previous few properties of Hilbert spaces are inherited from general inner product spaces. We now look at the Bessel inequality, which is exclusive to Hilbert spaces. This is useful as an auxilliary tool to prove that infinite linear combinations with coefficients given by inner products converge.

Theorem 2.6 (Bessel inequality)**.** *Let (e_j) be an orthonormal sequence in a Hilbert space H. Then for every $x \in H$,*

$$\sum_{j=1}^{\infty} |\langle x, e_j\rangle|^2 \le \|x\|^2.$$

Proof. The proof of the Bessel inequality requires a lot of algebra, so we prove it here with some algebra omitted since it is a frequently used inequality. This proof is adapted from [Con19].

We will start by showing that a few equations hold for finite orthonormal sets $\{e_1, \dots, e_n\}$ using inner product axioms and identities defined earlier. Then, we'll derive the Bessel inequality for finite sets. Lastly, we'll use a limiting process to show that the Bessel inequality holds for the infinite case of an orthonormal sequence. The first such equality is that

$$\left\langle x, \sum_{j=1}^{n} \langle x, e_j\rangle\, e_j \right\rangle = \sum_{j=1}^{n} \langle x, \langle x, e_j\rangle e_j\rangle = \sum_{j=1}^{n} |\langle x, e_j\rangle|^2,$$

which the reader can verify. Thus, skipping through the algebra, we have that

$$\begin{aligned}
\left\| x - \sum_{k=1}^{n} \langle x, e_k\rangle\, e_k \right\|^2 &= \left\langle x - \sum_{k=1}^{n} \langle x, e_k\rangle\, e_k, x - \sum_{j=1}^{n} \langle x, e_j\rangle\, e_j \right\rangle \\
&= \|x\|^2 - \left\langle x, \sum_{j=1}^{n} \langle x, e_j\rangle\, e_j \right\rangle - \overline{\left\langle x, \sum_{j=1}^{n} \langle x, e_j\rangle\, e_j \right\rangle} + \left\| \sum_{k=1}^{n} \langle x, e_k\rangle\, e_k \right\|^2 \\
&= \|x\|^2 - 2\sum_{j=1}^{n} |\langle x, e_j\rangle|^2 + \sum_{k=1}^{n} |\langle x, e_k\rangle|^2 \\
&= \|x\|^2 - \sum_{k=1}^{n} |\langle x, e_k\rangle|^2.
\end{aligned}$$

Now, given that we are looking at a norm in this second equation, we have that

$$0 \le \left\| x - \sum_{k=1}^{n} \langle x, e_k\rangle\, e_k \right\|^2 = \|x\|^2 - \sum_{k=1}^{n} |\langle x, e_k\rangle|^2,$$

so consequently we have the Bessel inequality for finite sets:

$$\sum_{k=1}^{n} |\langle x, e_k\rangle|^2 \le \|x\|^2.$$

Since for each k, we have that

$$|\langle x, e_k\rangle|^2 \ge 0,$$

we have that the sequence (y_n), where

$$y_n = \sum_{k=1}^{n} |\langle x, e_k\rangle|^2,$$

is bounded and increasing. Thus from the monotone convergence theorem, we have that (y_n) converges, say to some $y \in \mathbb{R}$. Since $y_n \le \|x\|^2$ for each n, we have that $\|x\|^2 \ge \lim y_n = y$. In other words,

$$\|x\|^2 \ge \sum_{k=1}^{\infty} |\langle x, e_k\rangle|^2.$$

■

The completeness of Hilbert spaces also allows us to reason about concepts that require preciseness, like orthogonality. Although we can define orthogonality for any inner product space, the completeness of a Hilbert space is required for such constructions as the orthogonal complement to be sensible.

Orthogonal Complement

We now build up a definition of the orthogonal complement, which will be key in proving the spectral theorem later in the essay. We'll apply the results from the previous section to Hilbert spaces to do so.

First, we prove that if a set is a convex subset of an inner product space, then for any point in the inner product space there is a unique 'closest' vector to that point in the subset. This result holds for any convex subset of the space, which need not be a vector subspace.

Theorem 2.7. *Let H be a Hilbert space and M a nonempty closed convex subset of H. Then for every given $x \in H$ there exists a unique $y \in M$ whose distance from x is minimized to a value $\delta \in \mathbb{R}$. In other words,*

$$\delta = \inf_{\tilde{y} \in M} \|x - \tilde{y}\| = \|x - y\|.$$

Proof. We start with the existence. Note that we can pick an arbitrary point $\tilde{y} \in M$ and see that $\|x - \tilde{y}\| \leq \inf_{\tilde{y} \in M} \|x - \tilde{y}\| < \infty$. By the definition of an infimum there is a sequence (y_n) in M such that we have the decreasing sequence

$$(\delta_n) \longrightarrow \delta, \qquad \delta_n = \|x - y_n\|.$$

We show that (y_n) is Cauchy. Writing $y_n - x = v_n$, we get $\|v_n\| = \delta_n$ and

$$\|v_n + v_m\| = \|y_n + y_m - 2x\| = 2\|\tfrac{1}{2}(y_n + y_m) - x\| \geq 2\delta.$$

Because M is convex, we have that $\frac{1}{2}(y_n + y_m) \in M$. Furthermore, we have $y_n - y_m = v_n - v_m$. By the parallelogram equality,

$$\begin{aligned} \|y_n - y_m\|^2 = \|v_n - v_m\|^2 &= -\|v_n + v_m\|^2 + 2(\|v_n\|^2 + \|v_m\|^2) \\ &\leq -(2\delta)^2 + 2(\delta_n^2 + \delta_m^2) \longrightarrow 0, \end{aligned}$$

so (y_n) is Cauchy. Since M is closed in X, we have a $y \in M$ such that $(y_n) \longrightarrow y$. Since $y \in M$, $\|x - y\| \geq \delta$. Furthermore, $\|x - y\| = \delta$ since

$$\|x - y\| \leq \|x - y_n\| + \|y_n - y\| = \delta_n + \|y_n - y\| \longrightarrow \delta.$$

Next, we show uniqueness. Suppose that $y \in M$ and $y_0 \in M$ both satisfy

$$\|x - y\| = \delta, \qquad \|x - y_0\| = \delta.$$

By the parallelogram equality,

$$\begin{aligned} \|y - y_0\|^2 &= \|(y - x) - (y_0 - x)\|^2 \\ &= 2\|y - x\|^2 + 2\|y_0 - x\|^2 - \|(y - x) + (y_0 - x)\|^2 \\ &= 4\delta^2 - 4\|\tfrac{1}{2}(y + y_0) - x\|^2. \end{aligned}$$

Notice that $\frac{1}{2}(y + y_0) \in M$, so

$$\|\tfrac{1}{2}(y + y_0) - x\| \geq \delta.$$

Thus we have that

$$0 \leq \|y - y_0\|^2 = 4\delta^2 - 4\|\tfrac{1}{2}(y + y_0) - x\|^2$$
$$\leq 4\delta^2 - 4\delta^2 = 0,$$

so $\|y - y_0\| = 0$ and clearly $y = y_0$. ■

We've now proven the existence of closest vectors to convex subsets of inner product spaces. Since vector spaces are convex, we first show that the line from a point x to its closest point y on a vector space is orthogonal to this vector space. This result gives us a version of orthogonal projection from finite dimensions, where much of our theory was built up using the now-obsolete coordinate-wise representations of vectors.

Lemma 2.8. *Let H be a Hilbert space containing the point x. For any closed subspace Y of H, if $y \in Y$ is the closest point to x then $z = x - y$ is orthogonal to Y.*

Proof. If we did not have that $z \perp Y$, then there would be $y_1 \in Y$ such that

$$\langle z, y_1 \rangle = \beta \neq 0.$$

Evidently, $y_1 \neq 0$. For any scalar α,

$$\begin{aligned}\|z - \alpha y_1\|^2 &= \langle z - \alpha y_1, z - \alpha y_1 \rangle \\ &= \langle z, z \rangle - \overline{\alpha}\langle z, y_1 \rangle - \alpha(\langle y_1, z \rangle - \overline{\alpha}\langle y_1, y_1 \rangle) \\ &= \langle z, z \rangle - \overline{\alpha}\beta - \alpha(\overline{\beta} - \overline{\alpha}\langle y_1, y_1 \rangle).\end{aligned}$$

Note that $\|z\| = \|x - y\| = \delta$ is the distance from x to Y, so taking

$$\overline{\alpha} = \frac{\overline{\beta}}{\langle y_1, y_1 \rangle}$$

yields

$$\begin{aligned}\|z - \alpha y_1\|^2 &= \langle z, z \rangle - \frac{\beta\overline{\beta}}{\langle y_1, y_1 \rangle} - \alpha\left(\overline{\beta} - \frac{\overline{\beta}}{\langle y_1, y_1 \rangle}\langle y_1, y_1 \rangle\right) \\ &= \langle z, z \rangle - \frac{|\beta|^2}{\langle y_1, y_1 \rangle} < \delta^2,\end{aligned}$$

which is a contradiction since Y is convex and therefore

$$\delta = \inf_{\tilde{y} \in Y} \|x - \tilde{y}\| < \|z - \alpha y_1\| = \|x - (y + \alpha y_1)\|.$$

Thus we must have that $z \perp Y$, so the lemma has been proven. ■

Now, recall the definition of the orthogonal complement.

Definition 2.9. Let X be an inner product space. The *orthogonal complement* of a vector subspace Y of X is the set

$$Y^{\perp} = \{z : \langle z, y \rangle = 0 \, \forall y \in Y\} \subseteq X.$$

In finite dimensions, we call this set the orthogonal complement because the inner product space X can be represented as the direct sum $Y \oplus Y^{\perp}$. We are now ready to show this in the case where Y is a closed subspace of a Hilbert space.

Theorem 2.10. *Let Y be any closed subspace of a Hilbert space H. Then*

$$H = Y \oplus Y^{\perp}.$$

Proof. Since H is complete and Y is closed, Y must be complete. Since Y is convex, by Theorem 2.7 and Lemma 2.8 we have that for every $x \in H$ there exists $y \in Y$ such that for some $z \in Y^{\perp}$,

$$x = y + z.$$

To prove uniqueness, suppose that $x = y + z = y_1 + z_1$, where $y_1 \in Y$, $z_1 \in Y^{\perp}$. Then $y - y_1 = z - z_1$, but since $y - y_1 \in Y$ while $z - z_1 \in Y^{\perp}$ we must have that $y - y_1 \in Y \cap Y^{\perp} = \{0\}$. Thus $y = y_1$, and by the same logic $z = z_1$. ■

This theorem allows us to more easily find a basis of our space. We start by finding a basis of a specific subspace and then finding a basis of its orthogonal complement, and then combine these two to get a full basis. This is the power that completeness has given us.

Separability

Let's return our attention to the existence of a Schauder basis. We'll make another definition here that will allow us to prove the existence of such bases.

Definition 2.11. A metric space M is said to be *separable* if it has a countable subset $\tilde{M}$ which is dense in M, i.e. the closure of $\tilde{M}$ is M.

Separability is nice because it allows us to approximate any member of our set with members of a countable subset. In particular, it can be shown that the separable Hilbert spaces are precisely those for which Schauder bases make sense.

Rather than discuss that point further, we'll jump straight ahead to orthonormal bases. Recall that a basis is orthonormal if $\langle e_i, e_j \rangle = \delta_{ij}$ for any two vectors. While we actually can't find orthonormal Hamel bases of separable spaces, because it can be shown they don't exist in such spaces, the added structure actually guarantees the existence of an orthonormal Schauder basis.

In fact, the neatest part about separable spaces is that we can use the familiar Gram-Schmidt orthogonalization process from finite dimensions in order to construct an orthonormal basis, much like we did in finite dimensions!

Theorem 2.12 (Gram-Schmidt)**.** *Every infinite-dimensional separable Hilbert space H has an orthonormal Schauder basis.*

Proof. Take a countable dense subset, i.e. one that can be arranged as a sequence (v_j) and which exists since our space is separable. We will orthonormalize (v_j) to a sequence (e_j) using the aforementioned Gram-Schmidt process while keeping the span of this sequence the same.

Take the first element of the sequence satisfying $v_j \neq 0$, and then set

$$e_1 = \frac{v_j}{\|v_j\|}.$$

Now, suppose that for the first n elements $v_1, \ldots, v_n$ we have found m nonzero orthonormal elements $e_1, \ldots, e_m$ where $m \leq n$ and

$$\operatorname{span}(e_1, \ldots, e_m) = \operatorname{span}(v_1, \ldots, v_n).$$

If v_{n+1} is in the span of $e_1, \ldots, e_m$, then our above equation holds for $v_1, \ldots, v_{n+1}$ in place of $v_1, \ldots, v_n$. Thus assume that $v_{n+1} \notin \operatorname{span}(e_1, \ldots, e_m)$. It follows that

$$w = v_{n+1} - \sum_{j=1}^{\infty} \langle v_{n+1}, e_j \rangle \, e_j \neq 0,$$

so that we can set

$$e_{m+1} = \frac{w}{\|w\|}.$$

Now notice that $e_{m+1} \perp \{e_1, \ldots, e_m\}$.

We can continue this process indefinitely, and ultimately we will either get an orthonormal set which may be finite or infinite and which we can arrange into a sequence (e_j). Now, we claim that any vector $u \in H$ that is perpendicular to each e_j is the zero vector. We prove this using the density of the elements v_n.

Choose u such that $u \perp e_j \, \forall j$. Since $\{v_j\}$ is dense in H, we have that there must be a sequence (w_k) of elements of $\{v_j\}$ with $(w_k) \longrightarrow u$. Each w_k is a finite linear combination of elements e_j by construction, so by the Bessel inequality

$$\|w_k\|^2 = \sum_{j \in \mathbb{N}} |\langle w_k, e_j \rangle|^2 = \sum_{j \in \mathbb{N}} |\langle u - w_k, e_j \rangle|^2 \leq \|u - w_k\|^2,$$

which we get because of the fact that $\langle u, e_j \rangle = 0 \, \forall j$. Thus we see that evidently $\|w_k\| \longrightarrow 0$, so $u = 0$.

We have now constructed a sequence (e_j) whose elements are orthonormal and linearly independent. Furthermore, the orthogonal complement of their Schauder span is $\{0\}$. This means that

$$H = \operatorname{span}(e_j) \oplus \{0\}$$

by 2.10, and therefore (e_j) span and consequently constitute an orthonormal Schauder basis. ■

The notion of a separable Hilbert space can be shown to be even stronger: just as every finite n-dimensional real inner product space is isomorphic to $\mathbb{R}^n$, we can show that any infinite-dimensional separable Hilbert space is isomorphic to the sequence space ℓ^2.

For our purposes, all that matters is that we've found a strong enough restriction of infinite-dimensional inner product spaces to state and prove a spectral theorem, once we've built up some more operator theory.

Some Background on Operators

Equipped with the orthogonal complement and separability, we are almost ready to prove a sufficiently strong spectral theorem for our use case. Up to now, we've been building up theory about the kinds of vector spaces on which we want to work, restricting to Hilbert spaces and Schauder bases. In this section, we'll build up some restrictions on the kinds of linear operators where our spectral theorem should hold.

We first recall the definition of an adjoint operator from finite-dimensional inner product spaces.

Definition 2.13. Let X be an inner product space and A a linear operator on X. The *adjoint* operator A^* is the unique linear operator such that $\langle w, Av \rangle = \langle A^* w, v \rangle$ for any $v, w \in X$. We call A *self-adjoint* if $A^* = A$.

Remember that for the finite-dimensional spectral theorem, we only required an operator to be self-adjoint. In infinite dimensions, this turns out to be insufficient. An intuitive reason why we need a stronger condition is the way that we chose unit eigenvectors in the finite-dimensional case. Implicitly, we used the fact that the unit sphere is compact in finite dimensions to choose these eigenvectors. The unfortunate reality of infinite-dimensional spaces is that the unit sphere is no longer compact. Thankfully, we can fix this problem by restricting the class of operators for which the spectral theorem holds a bit further.

We will now define a class of *bounded* operators. These operators will give us some of the structure needed to pick an eigenvector in an analogous fashion to finite dimensions.

Definition 2.14. Let X be an inner product space and A a linear operator. The operator A is said to be *bounded* if there is a real number c such that for all $x \in X$,

$$\|Ax\| \leq c\|x\|.$$

We'll now define an operator norm, which will allow us to quantify how much operators scale their inputs. We'll see that this quantification is crucial for finding a way around the problem of non-compactness of the unit ball.

Definition 2.15. Given a bounded operator A on a space X, we define its operator norm by

$$\|A\| = \sup_{x \in X \setminus \{0\}} \frac{\|Ax\|}{\|x\|}.$$

If $X = \{0\}$, we set $\|A\| = 0$.

The operator norm is the first step towards finding the right conditions for operators where the spectral theorem holds. What's nice is that we can restrict our search to the unit ball in the following way:

Lemma 2.16. *Let $A : H \to H$ be a bounded self-adjoint operator on a Hilbert space. Then*

$$\|A\| = \sup_{\|x\| \leq 1} |\langle A, x \rangle|.$$

As a reminder, since A is self-adjoint, $\langle Ax, x \rangle \in \mathbb{R}$.

Proof. For notational convenience, let

$$s(A) = \sup_{\|x\| \leq 1} |\langle A, x \rangle|.$$

By Cauchy-Schwarz, if $\|x\| \leq 1$ then

$$|\langle A, x \rangle| \leq \|Ax\| \|x\| \leq \|A\| \|x\|^2 \leq \|A\|,$$

so that $s(A) \leq \|A\|$.

The other side of the proof requires a more involved geometric argument. If $\lambda > 0$, then

$$\left\langle A\left(\lambda x \pm \frac{1}{\lambda} Ax\right), \lambda x \pm \frac{1}{\lambda} Ax \right\rangle = \langle A(\lambda x), \lambda x \rangle + \langle A^2(\tfrac{1}{\lambda} x), A(\tfrac{1}{\lambda} x) \rangle \pm 2\|Ax\|^2.$$

We can now subtract these equations from each other to get that

$$\begin{aligned} 4\|Ax\|^2 &= \left\langle A\left(\lambda x + \frac{1}{\lambda} Ax\right), \lambda x + \frac{1}{\lambda} Ax \right\rangle - \left\langle A\left(\lambda x - \frac{1}{\lambda} Ax\right), \lambda x - \frac{1}{\lambda} Ax \right\rangle \\ &\leq s(A)(\|\lambda x + \tfrac{1}{\lambda} Ax\|^2 + \|\lambda x - \tfrac{1}{\lambda} Ax\|^2) \end{aligned}$$

since the above inner products are of form $\langle Au, u\rangle$ and hence satisfy $|\langle Au, u\rangle| = |\langle A\hat{u}, \hat{u}\rangle| \|u\|^2 \leq s(A)\|u\|^2$. We can now apply the parallelogram equality to get that

$$4\|Ax\|^2 \leq 2s(A)(\lambda^2\|x\|^2 + \tfrac{1}{\lambda^2}\|Ax\|^2.$$

Assuming that $\|Ax\| \neq 0$, we can set $\lambda^2 = \frac{\|Ax\|}{\|x\|}$ to get that

$$4\|Ax\|^2 \leq 2s(A)\left(\frac{\|Ax\|}{\|x\|}\|x\|^2 + \frac{\|x\|}{\|Ax\|}\|Ax\|^2\right) = 4s(A)\|Ax\|\|x\|.$$

Hence for any $x \in H$, $\|Ax\| \leq s(A)\|x\|$, i.e. $\|A\| \leq s(A)$. ■

We'll now define a class of *compact* operators which will allow us to make a similar selection of unit eigenvectors when proving the spectral theorem in infinite dimensions.

Definition 2.17. Let V be an inner product space space and A a linear operator on V. We call A a *compact operator* if the closure of $A(B(0,1))$, the image of the unit ball, is compact.

In fact, compact operators turn out to be exactly the kind of operator we need to prove the spectral theorem. So, we finally turn our attention there.

Spectral Theorem

We are now in a position to state and prove a version of the spectral theorem that is sufficiently strong for our use case, namely to analyze the infinite-dimensional solution space of the wave equation.

Theorem 2.18 (Spectral Theorem for Compact Self-Adjoint Operators)**.** *Let H be a separable infinite-dimensional Hilbert space, and let A be a compact self-adjoint operator on H. Then there exists a sequence of real eigenvalues (λ_n) with $\lambda_n \to 0$ as $n \to \infty$, and a orthonormal Schauder basis v_n of eigenvectors with $Av_n = \lambda_n v_n$ for all $n \geq 1$.*

Remarkably, the proof of our more general spectral theorem follows the same pattern as the proof in finite dimensions, where for a given self-adjoint transformation we used the fact that it must have an eigenvector, which along with induction on the orthogonal complement of the span of this eigenvector gave us an orthonormal basis of our space.

We build up this proof via a few key lemmas. First, we recall the A^*-invariance of orthogonal complements of A-invariant subspaces, i.e. subspaces V such that $A(V) \subseteq V$.

Lemma 2.19. *Let $A : H \to H$ be a bounded operator on a Hilbert space. If $V \subseteq H$ is an A-invariant subspace, then $V^\perp$ is A^*-invariant.*

Proof. If $w \in V^\perp$ and $v \in V$ then $\langle A^*w, v\rangle = \langle w, Av\rangle = 0$. As $v \in V$ was arbitrary, we must have that $A^*w \in V^\perp$. ■

Just as in finite dimensions, this theorem will enable us to pick an eigenvector, whose span is by definition invariant, and reduce proof of the theorem to the orthogonal complement of this eigenvector. To make sure we can pick the unit eigenvector with maximal eigenvalue, we need to make sure that the infinite-dimensional unit sphere indeed has such a vector. Our next lemma proves this, finding the first eigenvector.

Lemma 2.20. *Let A be a compact self-adjoint operator on a non-trivial Hilbert space. Then at least one of $\pm\|A\|$ must be an eigenvalue of A.*

Proof. If $\|A\| = 0$ then $A = 0$ and we're done, so we may assume that $\|A\| > 0$. By lemma 2.16 we know that there is a scalar α with $|\alpha| = \|A\|$ and a sequence (x_n) in H such that for any $n \geq 1$, $\|x_n\| = 1$ where $\langle Ax_n, x_n\rangle \to \alpha$ as $n \to \infty$. Since we know that $\langle Ax_n, x_n\rangle$ is real, we know that $\alpha \in \{\pm\|A\|\}$. Now, we can note that since x_n are unit vectors,

$$\begin{aligned} 0 \leq \|Ax_n - \alpha x_n\|^2 &= \|Ax_n\|^2 - 2\alpha\langle Ax_n, x_n\rangle + \alpha^2 \\ &\leq 2\|A\|^2 - 2\alpha\langle Ax_n, x_n\rangle \to 2\|A\|^2 - 2\|A\|^2 = 0 \end{aligned}$$

as $n \to \infty$. What this means is that (Ax_n) converges iff (αx_n) converges, in which case they converge to the same limit.

As $\|x_n\| = 1$ is a bounded set and A is a compact operator, there is a subsequence (x_{n_k}) for which Ax_{n_k} converges, i.e. there is some $x \in H$ so that $Ax_{n_k} \to \alpha x$ as $k \to \infty$. By the above, we also have that $\alpha x_{n_k} \to \alpha x$, and so $x_{n_k} \to x$, as $k \to \infty$. So, $Ax = \alpha x$, and since x is the limit of unit vectors we find that x is therefore a unit eigenvector. ■

We can now combine the previous two lemmas to prove our spectral theorem for compact self-adjoint operators.

Proof of the Spectral Theorem. Our assumptions are that H is an infinite-dimensional Hilbert space with $A : H \to H$ a compact self-adjoint operator. We proceed by induction.

By lemma 2.20 we know that there exists an eigenvector e_1 with eigenvalue $\lambda_1 \in \mathbb{R}$, where $|\lambda_1| = \|A\|$. We can now also assume without loss of generality that $\|e_1\| = 1$. This concludes the base case.

Now, for the inductive case. Suppose for induction that we have already found orthonormal eigenvectors $e_1, \ldots, e_n$ with eigenvectors $\lambda_1, \ldots, \lambda_n$. Let $V_n = \langle e_1, \ldots, e_n\rangle$ be the span of these vectors, and note that $A(V_n) \subseteq V_n$ since these are eigenvectors. We can now apply lemma 2.19 to get that $A(V_n^\perp) \subseteq V_n^\perp$ since A is self-adjoint. Let's write the restriction of A to $V_n^\perp$ as

$$A_n = A|_{V_n^\perp} : V_n^\perp \to V_n^\perp.$$

Now, A_n inherits the properties of compactness and self-adjointness from A. Therefore, we can apply lemma 2.20 again to A_n to find another unit eigenvector e_{n+1} with eigenvalue λ_{n+1} where now $|\lambda_{n+1}| = \|A_{n+1}\|$. Thus we have an orthonormal set of $n+1$ eigenvectors $e_1, \ldots, e_{n+1}$, completing the inductive step.

Ultimately, we are left with an orthonormal sequence (e_n) of eigenvectors with $Ae_n = \lambda_n e_n$ and $\lambda_n \in \mathbb{R}$. We next show that $\lambda_n \to 0$ as $n \to \infty$. By construction,

$$|\lambda_{n+1}| = \|A_n\| = \left\|A|_{V_n^\perp}\right\| \leq \left\|A|_{V_{n-1}^\perp}\right\| = \|A_{n-1}\| = |\lambda_n|,$$

so that by induction

$$|\lambda_1| \geq |\lambda_2| \geq \cdots$$

and hence (λ_n) is a nonnegative decreasing sequence of reals.

Suppose for contradiction that this sequence does not approach 0 in the limit. Then, for some $\varepsilon > 0$, we have that $|\lambda_n| > \varepsilon \, \forall n \geq 1$. But then $\varepsilon e_n = A\left(\frac{\varepsilon}{\lambda_n} e_n\right) \in A(B(0,1))$ for all $n \geq 1$, and because $e_n \perp e_m$ for $n \neq m$, $\|\varepsilon e_n - \varepsilon e_m\| = \varepsilon\sqrt{2}$ for such vectors. This means that the sequence (εe_n) is contained in the closure of $A(B(0,1))$. However, the closure of $A(B(0,1))$ is compact, and so the fact that (εe_n) cannot have a convergent sequence is a contradiction.

Letting $V = \bigcup_n V_n = \langle e_1, e_2, \ldots \rangle$, we can also see that $A|_{V_n^\perp} = 0$ since

$$|\lambda_n| = \left\| A|_{V_n^\perp} \right\| \geq \left\| A|_{V^\perp} \right\|.$$

It just remains to construct our orthonormal basis of eigenvectors. Since $V^\perp \subseteq \ker A$ is a separable Hilbert space, either $V^\perp = \{0\}$ in which case we're done, or else $V^\perp$ must have a orthonormal Schauder basis by theorem 2.12, which we'll call (u_n). We now have an orthonormal Schauder basis (e_n) of A-eigenvectors of V, and an orthonormal Schauder basis of A-eigenvectors of $V^\perp$. By theorem 2.10, $H = V \oplus V^\perp$, and so taking the union of these two bases gives us a basis for H. Hence we've finished the proof. ■

Now armed with the spectral theorem, we are ready to turn back to our original problem. But, as the apt reader may have noticed, we need a bit of measure and distribution theory to apply the spectral theorem to our wave equation. Once we do that, we'll be set!

Measure and Distribution Theory

So, the last few sections discussed Hilbert spaces. The apt reader will have noticed that in fact $C^2(\Omega)$ is *not* a Hilbert space, because it's not complete. Alas, I was hoping to avoid a discussion of measure and distribution theory, but in the next few subsections we'll introduce the formally correct spaces for PDE solutions to live in, and then show how we can apply the spectral theorem to our case.

For the rest of this chapter, the definition of Lebesgue measure is assumed. We will go quickly through elementary definitions in measure theory, and would advise a reader who has not yet seen these topics to look through [SS05] for more explanation on measure. We'll then take a distributional view on these functions to define the weak derivative, skipping technicalities and without explicitly discussing distributions. A proper treatment of the subject of distribution theory can be found in [Str03].

Why doesn't the space of continuous functions work for us? Well, this space is not complete in the norm induced by our inner product:

Example. Consider the sequence

$$f_n(x) = \tanh\left(nx + \frac{L}{2}\right),$$

which is a Cauchy sequence in $C^2([0, L])$. However, $f_n \to f$ defined as follows:

$$f(x) = \begin{cases} 0, & x < \frac{L}{2}, \\ \frac{1}{2}, & x = \frac{L}{2}, \\ 1, & x > \frac{L}{2}. \end{cases}$$

Evidently, $f \notin C^2([0, L])$, so that this space is incomplete.

The usual way to force completeness is to relax our definition of continuity to a definition of *almost everywhere* continuity. So, define

$$\mathcal{L}^2(\Omega) = \left\{ f : \int_\Omega f^2 < \infty \right\}$$

to be the space of all square-integrable functions equipped with the integral as the 'inner product' (and hence 'norm' [6] $\|f\|_2 = \sqrt{\int f^2}$).

This space is almost a vector space, but it does not have a unique zero. So let $\mathcal{N}(\Omega) = \{f : \|f\|_2 = 0\} \subseteq \mathcal{L}^2(\Omega)$ denote the elements of norm 0. We're finally ready to define our Hilbert space.

Definition 2.21. Given a set Ω, we define $L^2(\Omega) = \mathcal{L}^2(\Omega)/\mathcal{N}$ to be the quotient space of square integrable functions, modulo functions with integral zero, on Ω.

The L^2-space is indeed an inner product space, and it's even complete! Unfortunately, the quotient above meant that we have now lost the ability to even evaluate functions pointwise. So, we have no hope of computing local information exact enough to take a classical derivative. So, how do we rectify this? We move away from thinking about the local characterisation of the derivative. First, we can start to think about 'evaluating a local region' of a function by looking at the behaviour it exhibits when integrated against a suitable *test function*. To encode the locality of test functions, we require compact support, and for nice differentiability properties we require smoothness.

Think of the test function as being like a thermometer bulb. A thermometer bulb measures temperature not at a particular point, but in a local area around itself. So too do test functions give us information about how large a function is in a local area.

This way of looking at functions is a 'distributional' view. How does our new perspective allow us to take a derivative? Well, now we can take what's called a *weak derivative* via integration by parts, and get another L^2 function back.

Definition 2.22. Let $f \in L^2(\Omega)$. Its weak derivative $f' \in L^2(\Omega)$ is characterized by action

$$\int_\Omega f'\phi = -\int_\Omega f\phi'$$

on test functions ϕ.

The reason we didn't have to care about the other term in the integration by parts is that our definition of test functions (specifically, the compact support) means that this term is always 0. Note that the notation is a bit simplified and concerns about existence and well-definedness are not addressed here. The interested reader is pointed to the aforementioned text on distribution theory ([Str03]) for a more principled approach to defining weak derivatives.

That being said, our new notion of derivative now allows us to define Sobolev spaces, which will finally suffice as Hilbert spaces whose elements are functions whose derivatives are sufficiently nice.

Definition 2.23. The Sobolev space $H^k(\Omega) \subseteq L^2(\Omega)$ is the set of all L^2-functions on Ω whose first k weak derivatives are also in L^2. This space is equipped with the norm

$$\|f\|_{H^k(\Omega)} = \sum_{i=0}^{k} \left\|f^{(i)}\right\|_{L^2(\Omega)}.$$

As our wave equation is a second-order PDE, we will proceed working on the Sobolev space $H^2(\Omega)$. There's one last catch: to apply the spectral theorem, we actually need to

[6] Technically, this is a seminorm.

enforce the Dirichlet boundary conditions. This is because we'll need to use integration by parts, and once again it would be really nice for us if the non-integral term of the integration by parts had value 0.

How do we enforce this? Unfortunately, pointwise limiting is no longer meaningful, so we can't say that we only want functions whose values approach 0 as the inputs approach the boundary $\partial\Omega$. Thankfully, we can once again build this definition using compactly-supported functions.

Definition 2.24. The space H_0^1 is the closure of the infinitely differentiable functions compactly supported in Ω in $H^1(\Omega)$.

Essentially, in the limit, we will be left with functions that 'die off' as they approach the boundary, as desired. Since we still want to restrict our attention to H^2 functions on Ω, we'll proceed to the next section only considering functions in the space

$$\mathcal{H}(\Omega) = H^2(\Omega) \cap H_0^1(\Omega).$$

To recap, functions f in this space satisfy the following conditions:

(1) Their first two weak derivatives f' and f'' are square-integrable.
(2) They disappear on the boundary, in particular for any $g \in L^2(\Omega)$ we have that

$$\int_{\partial\Omega} fg = 0.$$

At long last, we have a suitable function space to apply the spectral theorem we built up earlier.

3. Eigenvalues All the Way Down

Having now built up an appropriate function space for us to work in, and a spectral theorem to match, we finally return to our wave equation.

The Inverse Laplacian

We need a suitable operator to work with, which will end up being the *inverse Laplacian* operator. For reasons that will soon become clear, though, we'll first choose the negative Laplacian $A = -\nabla^2$. We already know that the Laplacian is linear by the linearity of partial derivatives. It follows that the negative and inverse Laplacians would also be linear. We now show that the negative Laplacian operator A is self-adjoint:

$$\begin{aligned}\langle f, Ag\rangle - \langle Af, g\rangle &= \int_\Omega (g\nabla^2 f - f\nabla^2 g)\, dx \\ &= \int_{\partial\Omega} \left(g\frac{\partial f}{\partial n} - f\frac{\partial g}{\partial n} \right) dS,\end{aligned}$$

Where the second line is due to Green's second identity. Now, recall that we have chosen our space of functions $\mathcal{H}(\Omega)$ such that

$$\int_{\partial\Omega} f\frac{\partial g}{\partial n}\, dS = \int_{\partial\Omega} g\frac{\partial f}{\partial n}\, dS = 0.$$

It follows that $\langle f, Ag\rangle - \langle Af, g\rangle = 0$. In particular, we can now conclude that $\langle f, Ag\rangle = \langle Af, g\rangle$, i.e. that A is self-adjoint. Since we intend to apply the Spectral theorem to the inverse Laplace operator A^{-1}, the fact that A is self-adjoint means that

$$\langle f, A^{-1}g\rangle = \langle AA^{-1}f, A^{-1}g\rangle = \langle A^{-1}f, AA^{-1}g\rangle = \langle A^{-1}f, g\rangle$$

and hence A^{-1} is self-adjoint as well.

It can also be shown that the inverse Laplacian is compact, although for brevity we will not prove this here. The interested reader is directed to [Eva10] for more details.

One last technicality: since the negative Laplacian is an operator

$$A : \mathcal{H}(\Omega) \to L^2(\Omega),$$

the inverse Laplacian operator A^{-1} has domain $L^2(\Omega)$. We haven't yet shown that this L^2-space is separable.

In order to show the separability of the space $L^2(\Omega)$, consider the set of rectangles with countable endpoints. Denote the set of indicator functions on these rectangles by $\mathcal{I}$. Now, let V denote the linear span[7] over $\mathbb{Q}$ of the indicator functions in $\mathcal{I}$. This set, containing the elements

$$\sum_{v_i \in I} q_i v_i$$

for finite subsets $I \subset \mathcal{I}$ and $q_i \in \mathbb{Q}$, is still countable. It can also be shown that it is dense in $L^2(\Omega)$. Hence our space $L^2(\Omega)$ is suitable to apply our spectral theorem.

We now have a compact self-adjoint operator A^{-1} on the separable Hilbert space $L^2(\Omega)$. We can finally think about the orthonormal Schauder eigenbasis (v_n) that the spectral theorem tells us will exist in this situation. As we'll see, we've already found this: it's our set of solutions from chapter 1.

The Spectrum of the Wave Equation

Recall that in the first chapter, we found a set of solutions to the wave equation under given conditions in both one and two-dimensional cases. We'll now show that these functions are the eigenfunctions of the inverse Laplacian, and hence by the spectral theorem any function describing the movement of our sound-producing media can be expressed as a linear combination of these functions.

When we solved the wave equations by separation of variables, we got a function u that was a product of a sinusoid in t, say with period ω, with some function w in the spatial variables. Now, this means that $\partial_{tt}\, u = \omega^2 u$, such that in both cases we see that for our function u the wave equation reduces to

$$\omega^2 u - \nabla^2 u = 0.$$

Hence, any solution to the wave equation with the time component separated is an eigenfunction of the Laplacian operator ∇^2 with eigenvalue ω^2. But by linearity, this means that such a function is also an eigenfunction of the negative Laplacian A, and in turn the negative Laplacian A^{-1}. Since we have found all such eigenfunctions[8] in both one and two dimensions in chapter 1, by the spectral theorem these eigenfunctions (after normalisation) constitute

[7] Here, we mean a linear span in the Hamel sense, i.e. that vectors must be represented as finite linear combinations of basis elements.

[8] In two dimensions, the polar separability of eigenfunctions is required by the wave equation's boundary conditions.

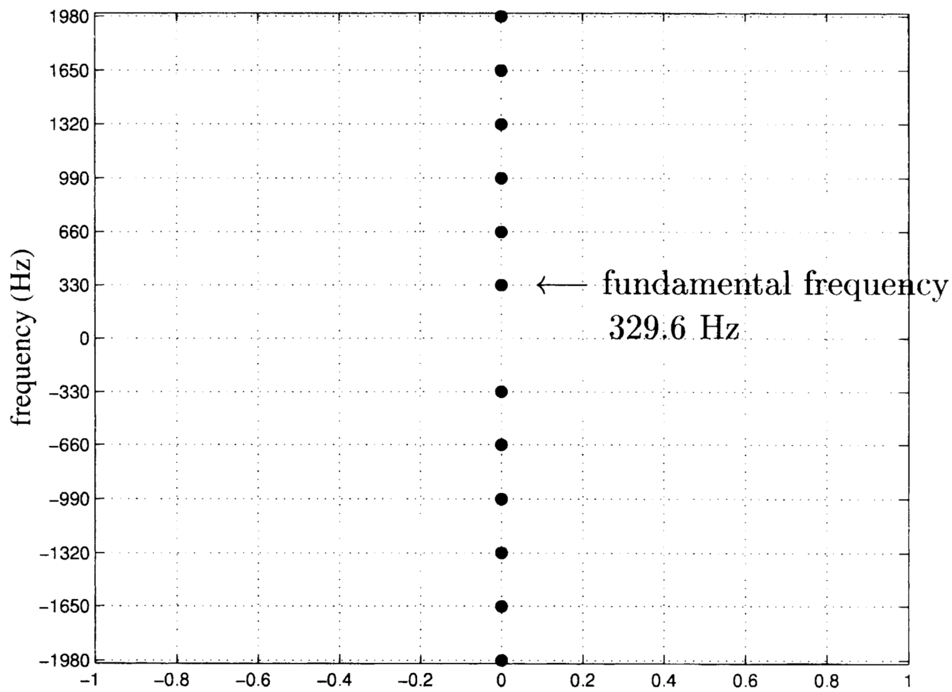

Figure 4. Frequencies of an ideal high E guitar string (image credit [HT01])

an orthonormal basis of $L^2(\Omega)$. Hence the spectrum, or set of eigenvalues, of ∇^2 is the set of eigenvalues of the solutions we found in chapter 1.

Now, since any solution to the wave equation is in $L^2(\Omega)$ by assumption, and the eigenfunctions form a basis of this space, it follows that the solutions of the wave equation can be expressed as an infinite linear combination of eigenfunctions. As an aside, one can also use this fact to solve the wave equation without explicitly considering initial conditions, by composing a solution as a suitable sum of eigenfunctions of the boundary valye problem. The coefficients of this sum are found via inner products of initial conditions and eigenfunctions.

Back to our main topic. As each eigenfunction has as a factor a sinusoid in t with fixed period, the frequencies emitted by the instruments correspond to the periods of the eigenfunctions. In particular, the frequencies are the square roots ω of the Laplacian eigenvalues. Therefore, no matter how each instrument is struck, the set of possible resonant frequencies stay the same. This is what the spectral theorem has bought us.

Let's think about the periods of these sinusoids. For the one-dimensional wave, the sinusoids have periods $\frac{cn\pi}{L}$, and these frequencies are in a 'Harmonic series' [9]. We see this in Figure 4, which shows the frequencies of a guitar. Note that the case of a violin is similar, but as the violin's E string is tuned one octave higher the frequencies would all need to be doubled.

On the other hand, on a timpani these frequencies would be $cj_{n,m}$, where $j_{n,m}$ are the zeros of the Bessel function. These are not in Harmonic series, as can be seen in Figure 5.

[9]This is called a Harmonic series in music terms; in mathematics, we'd call this a geometric progression.

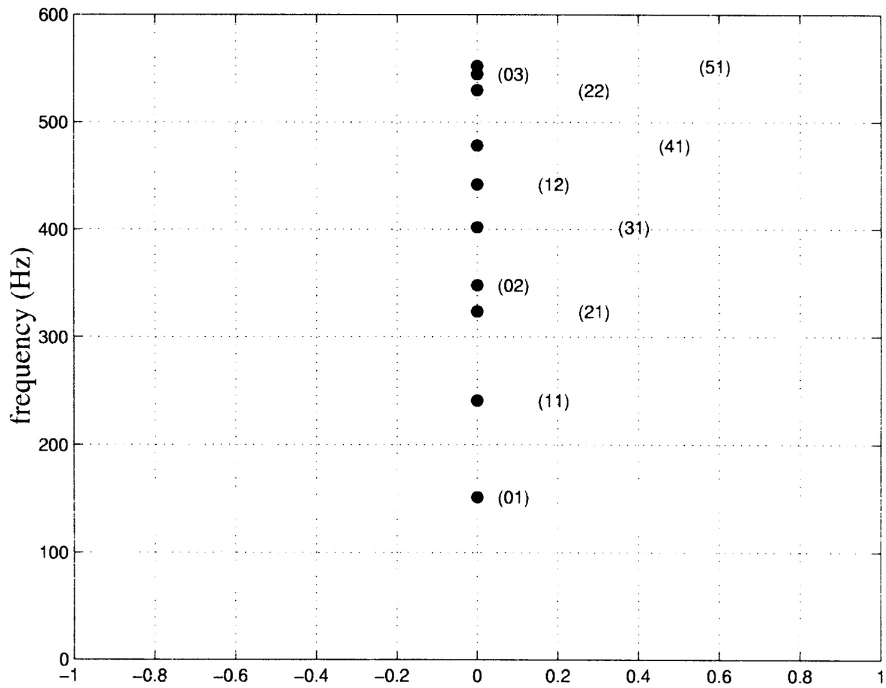

Figure 5. Frequencies of an ideal drum (image credit [HT01])

Resonant Frequencies of the Timpani

So, why does the timpani sound pleasing if the frequencies we've calculated it to produce are so un-musical? Surprisingly, as can be seen in Figure 6, the real-world frequencies produced by a timpani do not match these zeros of the Bessel function! Note that each frequency here also has an associated decay rate, which indicates how quickly the sound produced by the eigenfunction corresponding to a certain frequency dies out.

According to [FR13], the main real-world phenomenon accounting for this difference is called 'air loading', meaning that the timpani moves such a great mass of air that the oscillation of the drum membrane becomes coupled with that of the air. While we won't go into the mathematics of this coupling, we will make two key observations here about why this effect makes the timpani sound musical.

First, the frequencies of the timpani labeled $(n1)$ are pushed closer to being in Harmonic series. Originally, these frequencies corresponded to the zeros $j_{n,1}$ of the Bessel function J_1.

Second, the eigenfunctions associated with these frequencies, called $(n1)$ modes, decay due to energy loss much more slowly than the other frequencies, meaning that they are more prominent for a listener. As a result, for $n > 0$, the $(n1)$ modes are called the *preferred modes* of the timpani. So, rather than the fundamental frequency of the timpani produced by the (01) mode dominating, it is instead the 'preferred frequency' produced by the (11) mode that is most audible when a timpani is played.

The reason why the (11) mode dominates also has to do with how the timpani is struck, since striking it in a way that excites this mode [10] produces more clean acoustics than striking

[10] In particular, it is struck such that the initial conditions have high inner product with the (11) mode.

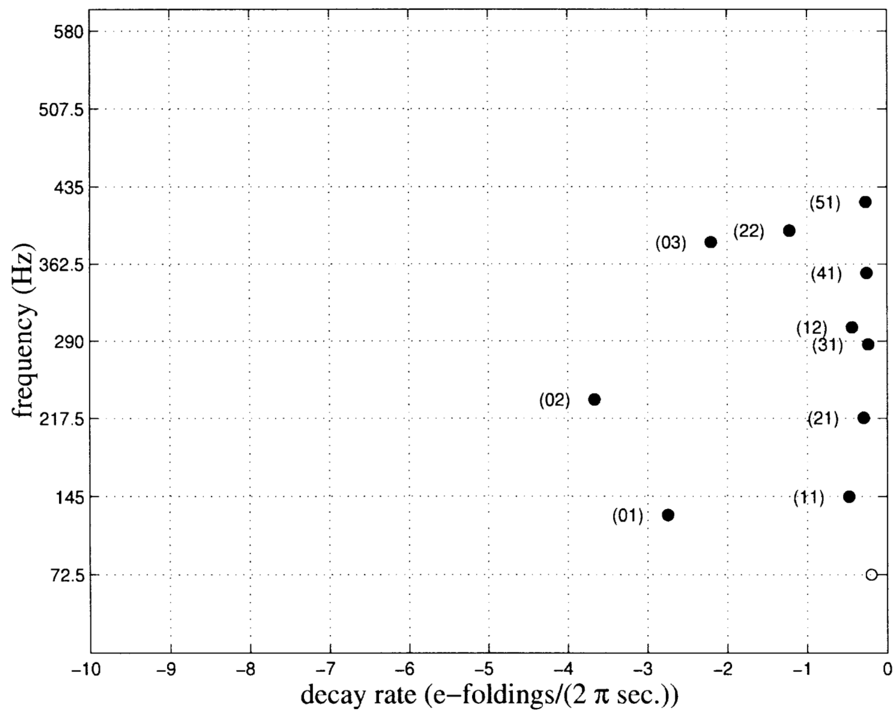

Figure 6. Frequencies and decay rates of a timpani as measured experimentally in the real world (image credit [HT01])

the (01) mode. In any case, this analysis is sufficient to answer our original questions from the introduction.

Why is a Timpani So Hard to Tune?

After a very long mathematical journey, we're finally ready to discuss the reason why my band director didn't let us tune the timpanis all those years ago.

Let's start with the most obvious reason, which has to be our assumption of uniform tension. In order for a our mathematical analyses to be valid, both timpanis and violins need uniform tension on the boundary of their sound-producing medium; i.e., their heads and strings, respectively.

For a violin, the boundary is two points, one of which is fixed at the bridge. So the only adjustment that a violinist must make in order to change the tension on a given part of the string is to turn the peg. The construction of the violin ensures that the string stays taut regardless of how the peg is tuned and hence the uniform tension assumption is maintained regardless.

On the timpani, however, one must tune six to eight lugs across the head of the drum to have the same tension. This requires a level of precision that was not needed for the case of a violin, but could be trained over time. In any case, this uniform tension assumption is much harder to achieve when tuning a timpani, whereas it is guaranteed when tuning a violin.

Anyone could have thought about uniform tension. Did our mathematics elucidate another potential reason for the timpani being harder to tune than the violin? Yes, and that's the fact that the eigenvalues of a timpani are not in harmonic series while those of a violin are. This means that a timpani relies on the physical mechanism of air loading to produce nice sound, and as discussed previously we are tuning to the frequency of the preferred (11) mode and not the fundamental.

However, the effects of these deviations from our mathematical theory are in fact quite predictable, so that an apt tuner can adjust their approach to account for these deviations. As a result, a tuner doesn't have to actively think about these complicated phenomena. So, ultimately, their effect on the difficulty of timpani tuning is not all that pronounced.

What makes a timpani so hard to tune? Frankly, it comes down to the fact that achieving equal tension across six lugs on the boundary requires the tuner to have precise ears and deft hands. As for why my band director didn't think any of the percussionists in the band were qualified to overcome these difficulties, I've got some conjectures. The amount of training it would take to tune the timpanis accurately would not have been worth the effort if the instrument was only tuned once every few months to a year. However, even if the instrument was tuned more frequently, I doubt that it would have produced a noticeable effect on the sound of a band at our level, despite what the Reddit experts say.

But in a situation like this where we don't know the answer, Occam's razor tells us that the simplest answer is most likely to be true. In this case, the simplest reason why my band director wouldn't have allowed us to tune the timpani is because teenagers are not known to exhibit the patience required for this challenging of a tuning.

Acknowledgements

The author would like to thank Andrew Dancer for his mathematical guidance and helpful resources for the project, as well as the wonderful friends Alice McNaughton-Scally, Eric Liang, Julian Xiao, and William Lacey for the conversations and feedback. Special thanks to Raphael dela Cruz and Ivy Wong for their help in editing the final draft of this essay.

References

[Axl97] Sheldon Jay Axler. *Linear Algebra Done Right*. Undergraduate Texts in Mathematics. Springer, New York, 1997. URL: http://linear.axler.net/.

[Bha22] Arav Bhattacharya. Linear algebra in infinite dimensions, 2022. URL: http://simonrs.com/eulercircle/irpw2022/arav-hilbert-paper.pdf.

[Car05] Neal L Carothers. *A short course on Banach space theory*. Number 64. Cambridge University Press, 2005.

[Con19] John B Conway. *A course in functional analysis*, volume 96. Springer, 2019.

[Eva10] Lawrence C. Evans. *Partial differential equations*. American Mathematical Society, Providence, R.I., 2010.

[EW18] Manfred Einsiedler and Thomas Ward. *Functional analysis, spectral theory, and applications*. Springer International, 2018.

[FR13] Neville Horner Fletcher and Thomas D. Rossing. *The Physics of Musical Instruments*. Springer International Publishing, 2013.

[Hew] Ian Hewitt. A6 differential equations ii. URL: https://courses.maths.ox.ac.uk/pluginfile.php/104991/mod_resource/content/5/Differential%20Equations%202.pdf.

[HT01] V.E Howle and Lloyd N Trefethen. Eigenvalues and musical instruments. *Journal of Computational and Applied Mathematics*, 135(1):23–40, Oct 2001. doi:10.1016/s0377-0427(00)00560-4.

[Jon21] Richard K. Jones. Chapter 2: Modes and nodes, Dec 2021. URL: https://wtt.pauken.org/.

[Kac66] Mark Kac. Can one hear the shape of a drum? *The American Mathematical Monthly*, 73(4):1, Apr 1966. doi:10.2307/2313748.

[Kre91] Erwin Kreyszig. *Introductory functional analysis with applications*, volume 17. John Wiley & Sons, 1991.

[Mel14] Richard Melrose. Introduction to functional analysis: 18.102 – spring 2014, 2014. URL: https://math.mit.edu/~rbm/18-102-S14/.

[mod] Preferred timpani modes. URL: http://www.hsc.edu.kw/student/materials/Physics/website/hyperphysics%20modified/hbase/music/timpani.html.

[Mor] David Morin. Waves (draft). URL: https://davidmorin.physics.fas.harvard.edu/books/waves/.

[Oli] Jim Oliver. M5 fourier series & pdes. URL: https://courses.maths.ox.ac.uk/pluginfile.php/104667/mod_resource/content/1/FSPDE-Lectures-Notes.pdf.

[Qua11] Quartl. *Illustration of the unit circles of some p-norms*. Nov 2011. URL: https://commons.wikimedia.org/w/index.php?curid=17428655.

[SS05] Elias M Stein and Rami Shakarchi. *Real analysis*. Princeton lectures in analysis. Princeton Univ. Press, Princeton, NJ, 2005. URL: https://cds.cern.ch/record/1385521.

[Str03] Robert S Strichartz. *A Guide To Distribution Theory And Fourier Transforms*. 2003.

Email address: bhattacharya.arav05@gmail.com