DS 4: Symmetries and Coupled Oscillators

Philip Cherian

February 11, 2022

1 Symmetry as an operation

The notion of symmetry is one that we are familiar with. For example, if I were to ask the following question: *which is more symmetric, a square or a circle?*, you would all probably intuitively know the answer. If I asked you to justify your answer, you would probably say something as follows:

(a) If we rotate a circle by any angle about its centre, it looks identical to itself. However, for a square, we need to rotate it by exactly 90 degrees in order for this to happen. As such, the circle has more "actions" that can keep it symmetric, as opposed to a square.

(b) Similarly, if I placed a mirror passing diametrically through the centre of the circle at any angle, then the half-circle and its reflection would look identical to the full circle. However, for a square, there are only four such axes of reflection which leave the figure unchanged.

These are all true. However, you might not be able to formulate this into a mathematically coherent form. Let us now try to make this more rigorous.

1.1 The symmetry matrix

To begin, let us consider the case in which we are dealing with discrete objects: say, masses on a beaded hoop, as shown in Figure (1). Now, imagine that every mass could be represented as components of a single vector. Then, if we start counting from the mass directly at the top and move clockwise, we can see that the total "state" of the system can be described by the vector

$$\underline{V} = \begin{pmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \\ \vdots \\ \theta_N \end{pmatrix}$$

Now, as we saw, the entire system would look identical if we rotated it (say) counter-clockwise by an angle $\phi = 2\pi/N$. In other words, if you happen to leave the room for a cup of coffee and one of your friends came by and rotated the system by this angle, you wouldn't be able to tell anything was amiss when you returned. Clearly, this is equivalent to relabelling our system, with the mass formerly known as θ_1 being labelled θ_2, and so on. Such a transformation is linear, since we can find a matrix that can transform the original vector to the new one. You should be able to do this relatively painlessly, since it's a matrix that – when acted on a vector – moves everything down by one step. As such, it is a matrix with all the next-to-diagonal entries being 1, one left-most corner entry being 1, and all the others being 0.

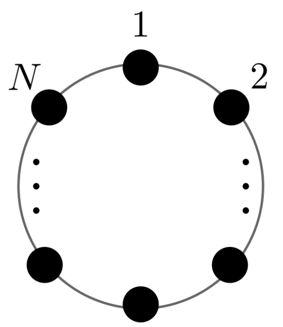

Figure 1: N masses are placed on a hoop in a symmetric fashion, separated by equal angles subtended at the centre of the hoop. The system is thus symmetric under rotations by any integral multiple of the angle $\phi = 2\pi/N$.

$$\underline{V}' = \begin{pmatrix} \theta_2 \\ \theta_3 \\ \vdots \\ \theta_N \\ \theta_1 \end{pmatrix} = \begin{pmatrix} 0 & 1 & 0 & 0 & \dots & 0 \\ 0 & 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \dots & \vdots \\ 1 & 0 & 0 & 0 & \dots & 0 \end{pmatrix} \cdot \begin{pmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \\ \vdots \\ \theta_N \end{pmatrix} = \underline{\underline{S}} \cdot \underline{V} \tag{1}$$

Of course, this is just the action of one such rotation. But it should be clear that we could just as well have rotated the system the angle $2 \times 2\pi/N$, and we'd *still* have no way of telling that the system has changed. We could, of course, write down a matrix for such a transformation, but we don't need to because we can just realise that such a transformation is essentially just two successive rotations by an angle $2\pi/N$. Thus, it should be obvious that

$$\underline{V}'' = \begin{pmatrix} \theta_3 \\ \theta_4 \\ \vdots \\ \theta_1 \\ \theta_2 \end{pmatrix} = \begin{pmatrix} 0 & 1 & 0 & 0 & \dots & 0 \\ 0 & 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \dots & \vdots \\ 1 & 0 & 0 & 0 & \dots & 0 \end{pmatrix} \cdot \begin{pmatrix} \theta_2 \\ \theta_3 \\ \vdots \\ \theta_N \\ \theta_1 \end{pmatrix} = \begin{pmatrix} 0 & 1 & 0 & 0 & \dots & 0 \\ 0 & 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \dots & \vdots \\ 1 & 0 & 0 & 0 & \dots & 0 \end{pmatrix}^2 \cdot \begin{pmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \\ \vdots \\ \theta_N \end{pmatrix} = \underline{\underline{S}}^2 \cdot \underline{V} \tag{2}$$

You should also be able to see that – because of the structure of space – if we repeated this process n times, we should come back to our starting point: this is not a property of the system, but of space itself, since by definition, a rotation by 2π leaves the system unchanged. The mathematical consequence of this is that the successive action of $\underline{\underline{S}}$ N times leaves the system unchanged. Therefore, such a transformation is the identity transformation. As a result, we have a powerful statement:

$$\underline{\underline{S}}^N = I_{N\times N} \tag{3}$$

Exercise: Explicitly show the above relation for the case where $N = 3$.

Exercise: Show that the determinant of $\underline{\underline{S}} = 1$.

It turns out that this has great consequences to the eigenvalues and eigenvectors of the matrix $\underline{\underline{S}}$, as we shall now see.

1.2 The eigenvalues of the symmetry matrix

Let us first establish a simple but powerful result: any power n of the matrix $\underline{\underline{S}}$ has the same eigenvectors as $\underline{\underline{S}}$, with eigenvalues being the original eigenvalues raised to n. It should be quite easy to show this: suppose that $\underline{v}$ is an eigenvector of $\underline{\underline{S}}$ with eigenvalue β. In this case,

$$\underline{\underline{S}}^n \cdot \underline{v} = \underline{\underline{S}}^{n-1}\underline{\underline{S}} \cdot \underline{v} = \beta \underline{\underline{S}}^{n-1} \cdot \underline{v} = \beta^2 \underline{\underline{S}}^{n-2} \cdot \underline{v} = \ldots = \beta^n \underline{v}. \tag{4}$$

Interestingly, I am not able to show trivially that *all* the eigenvectors of $\underline{\underline{S}}^n$ are the eigenvectors of $\underline{\underline{S}}$. Perhaps this is not true in general. If any of you can come up with a proof, I'd be interested to hear it.

Addendum: Thinking about it, it's clear that all the eigenvectors of $\underline{\underline{S}}^n$ need not be eigenvectors of $\underline{\underline{S}}$. The proof is simple by the contrapositive method: consider a matrix such that $\underline{\underline{S}}^n = I$ (like we have right now). Now, any arbitrary vector is an eigenvector of the identity I by definition. However, clearly not every vector is an eigenvector of $\underline{\underline{S}}$! Thus, not every eigenvector of $\underline{\underline{S}}^n$ is an eigenvector of $\underline{\underline{S}}$.

Because of the above result, we can very easily compute the eigenvalues of $\underline{\underline{S}}$ as follows. Suppose that $\underline{v}_k$ represents an eigenvector of $\underline{\underline{S}}$ with eigenvalue β_k. In this case, we know that

$$\underline{\underline{S}}^N \cdot \underline{v}_k = \beta_k^N \underline{v}_k \qquad \Longrightarrow \qquad \beta_k^N = 1, \tag{5}$$

where in the last step we've used the fact that the eigenvalues of the identity matrix are all 1. As a result, it should be clear that the eigenvalues β_k are the $N-$roots of unity!

1.3 The roots of unity

If you're not used to working with complex numbers, solving the above equation might seem very strange to you. The way this is usually done is the following: we write the right hand side (i.e. the number 1) as

$$1 = e^{i2\pi} = e^{i2\pi k}, \qquad k \in \mathbb{Z}, \tag{6}$$

which can easily be seen from the periodicity of the complex exponential. You could also use Euler's formula $\exp(i\theta) = \cos\theta + i\sin\theta$ to see this explicitly using the more familiar periodicity of the sine and cosine functions. Now, we see that because of the presence of k, when the root of 1 is taken with respect to N, we have

$$\beta_k = e^{i2\pi k/N}. \qquad k \in 0, 1, 2, \ldots N-1 \tag{7}$$

Notice how now there are only N allowed values of k, ranging from 0 to $N-1$. This should not be surprising, since the equation we're solving for is a polynomial equation of order N, and it should have at most N roots.

Exercise: Show that if $k = N$, you get back β_0, and if $k = 1$, you get back β_1, and so on.

The roots of unity have very interesting properties and I won't list all of them, but here are a couple that you might find useful. You should try to show all of these properties.

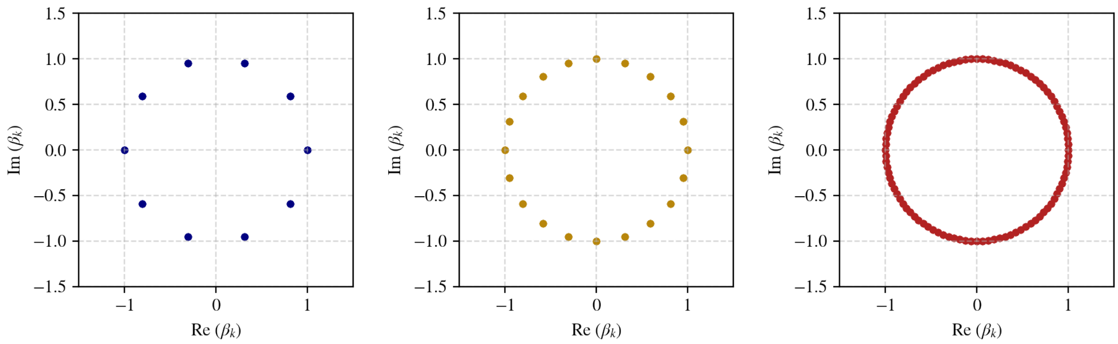

(a) $N = 10$ **(b)** $N = 20$ **(c)** $N = 100$

Figure 2: Plots of the N−th roots of unity for different values of N. The roots always lie on a unit circle, and as N increases, they become more and more dense.

(a) These are all complex numbers (except for the ± 1), and they all lie on a unit circle on the complex plane, separated by a constant angle $2\pi/N$, as shown in Figure (2).

(b) As N increases, these roots get squished together, but they still lie on the same unit circle.

(c) Multiplying each of these roots by $\beta_1 = e^{i2\pi/N}$, which is equivalent to a anticlockwise rotation in the complex plane by $2\pi/N$, moves you from one root to the other. Thus,

$$\beta_{k+1} = \beta_1 \beta_k \tag{8}$$

(d) If β is some complex root of unity, then so is β^{-1}. Furthermore, if $\beta \neq \pm 1$, then β^{-1} is a *different* root, since $\beta^{-1} = \beta^*$, the complex-conjugate of β. Thus, taking the inverse of β gives a reflection about the real axis.

Exercise: Show that when N is odd, $\beta = -1$ is not an N−th root of unity.

Exercise: Show that when N is even, there are 2 real and $N-2$ complex roots, and when N is odd there is only 1 real root and $N-1$ complex roots. In each case, there are an even number of complex roots of unity.

1.4 The eigenvectors of the symmetry matrix

Given that we've found the eigenvalues of $\underline{\underline{S}}$, we can easily find the eigenvectors. Suppose that $\underline{A}^k$ represents an eigenvector of $\underline{\underline{S}}$. Notice that since we've assumed that $\underline{A}$ is an eigenvector of $\underline{\underline{S}}$,

$$\underline{\underline{S}} \cdot \begin{pmatrix} A_1^k \\ A_2^k \\ A_3^k \\ \vdots \\ A_N^k \end{pmatrix} = \begin{pmatrix} A_2^k \\ A_3^k \\ \vdots \\ A_N^k \\ A_1^k \end{pmatrix} = \beta_k \begin{pmatrix} A_1^k \\ A_2^k \\ A_3^k \\ \vdots \\ A_N^k \end{pmatrix} \tag{9}$$

Thus, we can compare the entries on the left and right hand sides to say that

$$\begin{aligned} A_2^k &= \beta_k A_1^k \\ A_3^k &= \beta_k A_2^k = \beta_k^2 A_1^k \\ &\vdots \\ A_N^k &= \beta_k^{N-1} A_1^k \end{aligned} \tag{10}$$

Now, since we have the liberty to rescale an eigenvector, we can choose to set one of its component to be 1. If we choose $A_1^k = 1$, we have

$$\underline{A}^k = \begin{pmatrix} 1 \\ \beta_k \\ \beta_k^2 \\ \vdots \\ \beta_k^{N-1} \end{pmatrix} \tag{11}$$

Exercise: Using the explicit form of $\underline{A}^k$ given above, show that the action of $\underline{\underline{S}}$ (which corresponds just to shifting every element of the vector down by one "rung") corresponds to multiplying the entire vector by β_k.

Note: Remember that there are N different values of β_k here, one for each root of unity, each of which corresponds to a different eigenvector $\underline{A}^k$. The elements of each of the eigenvectors are the different powers of each β_k. You should be able to see that the elements of the different eigenvectors will all be the same, just permuted in different ways!

2 Discrete rotational symmetry

If a system is truly invariant under the action of symmetry, then the physics of a system should not depend on which of the different symmetric configurations we choose to work with with. For example, let us

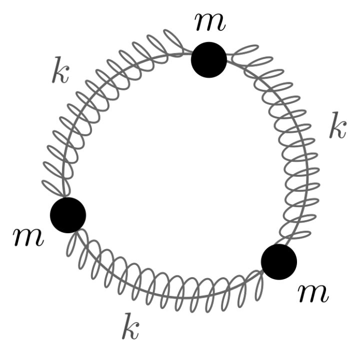

Figure 3: Consider a simple physical system represented above, with three identical masses connected by three identical springs.

consider a system with three masses connected by identical springs. You and two of your friends decide to solve this system, but you each decide to choose a different “first” mass. As a result, what you call mass 1 might be what your friend calls mass 2 or mass 3. However, since the system is truly symmetric about this choice, you should get exactly the same equations of motion for the system. In particular, suppose you choose the top-most mass to be mass 1. In this case, you would be working with a vector that looks like

$$\underline{V} = \begin{pmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \end{pmatrix}.$$

However, your friend would be working with a vector which would look like

$$\underline{V}' = \begin{pmatrix} \theta_2 \\ \theta_3 \\ \theta_1 \end{pmatrix}.$$

If you now solved for the complicated motion for the system, and got different equations for $\underline{V}(t)$ as a function of time, it would not in general be the same as the vector $\underline{V}'(t)$ that your friend obtained. Now, if your friend wishes to compare her results with yours, she would need a way to *transform* your results to be compatible with hers. If you think about this for a bit, you should realise that all she needs to do is act the symmetry matrix for this problem (the 3×3 symmetry matrix S), since we know that $\underline{V}' = \underline{\underline{S}} \cdot \underline{V}$.

We can, however, be even more restrictive: suppose you had obtained the following (matrix) equations for the motion of the three masses

$$\frac{\mathrm{d}\underline{V}}{\mathrm{d}t} = -\underline{\underline{\Lambda}} \cdot \underline{V}, \tag{12}$$

and your friend obtained

$$\frac{\mathrm{d}\underline{V}'}{\mathrm{d}t} = -\underline{\underline{\Lambda}}' \cdot \underline{V}'. \tag{13}$$

But we know that the choice of one of these symmetric configurations should not modify the underlying *physics*, meaning that the matrix $\underline{\underline{\Lambda}}$ which contains all of the physical content of the problem should be the same between you and your friend. In other words, we need that $\underline{\underline{\Lambda}} = \underline{\underline{\Lambda}}'$. This is a curious condition, so spend some time thinking about it. The idea is this: since all of the choices for $\underline{V}$ are equivalent, the matrix $\underline{\underline{\Lambda}}$ should not depend on which of the symmetric configurations you chose to define your system. Indeed, if you look at your system, it should be quite obvious that the symmetry indicates that both you and your friend have exactly the same $\underline{\underline{\Lambda}}$, since it only depends on the springs and masses on either side, which are identical!

Let us now see what the consequences of this are. Since the $\underline{\underline{S}}$ is a constant, we can take your equation and act $\underline{\underline{S}}$ on it to try and obtain the equation that your friend had. In this case:

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\underline{\underline{S}}\cdot\underline{V}\right) = -\underline{\underline{S}}\cdot\underline{\underline{\Lambda}}\cdot\underline{\underline{S}}^{-1}\left(\underline{\underline{S}}\cdot\underline{V}\right), \tag{14}$$

where we have used the fact that $I = \underline{\underline{S}}^{-1}\cdot\underline{\underline{S}}$. Now, since we've seen that $\underline{V}' = \underline{\underline{S}}\cdot\underline{V}$, we can compare this equation to the one given immediately above to show that

$$\underline{\underline{\Lambda}}' = \underline{\underline{S}}\cdot\underline{\underline{\Lambda}}\cdot\underline{\underline{S}}^{-1}. \tag{15}$$

Of course, our physical requirement of symmetry further enforces that $\underline{\underline{\Lambda}}' = \underline{\underline{\Lambda}}$. Putting all of this together, we get a very important result:

$$\underline{\underline{S}}\cdot\underline{\underline{\Lambda}} = \underline{\underline{\Lambda}}\cdot\underline{\underline{S}} \qquad\Longleftrightarrow\qquad \underline{\underline{S}}\cdot\underline{\underline{\Lambda}} - \underline{\underline{\Lambda}}\cdot\underline{\underline{S}} \equiv \left[\underline{\underline{\Lambda}},\underline{\underline{S}}\right] = 0. \tag{16}$$

When we see something like this, we say that the matrices *commute*, since it doesn't matter in which order we apply them to a vector. What this means is that if you act the matrix $\underline{\underline{\Lambda}}$ on your vector (which would give you your accelerations for the different masses), and then transformed it to your friend's coordinates by acting $\underline{\underline{S}}$ on it, it would give you exactly the same results as if you friend took her state vector, acting on it by $\underline{\underline{S}}$ to bring it to your coordinates and acted $\underline{\underline{\Lambda}}$ on it.

2.1 Commuting matrices

When you have two commuting matrices, it turns out that many of the properties of the matrices are shared. In particular, the fact that two matrices commute tells us that we can find a common eigenbasis. i.e., the eigenvectors of one of the matrices is the same as the eigenvectors of the other one.

You should now see why we went through all the trouble to find the eigenvectors of $\underline{\underline{S}}$. These vectors were much easier to compute than the eigenvectors of $\underline{\underline{\Lambda}}$, but – given that the two matrices commute – we can now be certain that the eigenvectors of $\underline{\underline{S}}$ are *also* eigenvectors of $\underline{\underline{\Lambda}}$.[1]

We can prove this quite simply in the case of non-degenerate eigenvalues. Suppose that all the eigenvectors of $\underline{\underline{S}}$ have distinct eigenvalues (as they do, which you should be able to see from our calculations in the previous section). The proof is as follows: begin by assuming that $\underline{A}$ is an eigenvector of $\underline{\underline{S}}$ with an eigenvalue β. In this case, we know that

$$\underline{\underline{S}}\cdot\underline{A} = \beta\,\underline{A}. \tag{17}$$

[1] There is a very important caveat here that you will appreciate during your course on Quantum Mechanics next semester. It is very easy to show that if $[P,Q] = 0$ for two finite-dimensional matrices, then every eigenvector of P is an eigenvector of Q and vice-versa, provided that the system has no degeneracy, i.e., every eigenvector has a distinct eigenvalue. If this is not the case, the result is *still* true in a restricted sense: not *every* eigenvector of P is an eigenvector of Q, but it is possible to form *one* eigenbasis which is common between the two. This involves a little more work to prove, and is usually done at the start of a Quantum Mechanics course.

We can now act this equation on both sides by $\underline{\underline{\Lambda}}$. In this case, since β is a constant, we can write:

$$\underline{\underline{\Lambda}} \cdot \underline{\underline{S}} \cdot \underline{A} = \beta \underline{\underline{\Lambda}} \cdot \underline{A}. \tag{18}$$

Now, since the matrices commute, we can write $\underline{\underline{\Lambda}} \cdot \underline{\underline{S}} = \underline{\underline{S}} \cdot \underline{\underline{\Lambda}}$, and therefore we have

$$\underline{\underline{S}} \cdot \left(\underline{\underline{\Lambda}} \cdot \underline{A}\right) = \beta \left(\underline{\underline{\Lambda}} \cdot \underline{A}\right). \tag{19}$$

In other words, if $\underline{A}$ is an eigenvector of $\underline{\underline{S}}$, then *so is* $\underline{\underline{\Lambda}} \cdot \underline{V}$! More importantly, $\underline{\underline{\Lambda}} \cdot \underline{V}$ is an eigenvalue of $\underline{\underline{S}}$ with exactly the same eigenvalue as $\underline{A}$! However, we claim that the eigenvectors of $\underline{\underline{S}}$ are non-degenerate: meaning that every eigenvector has only one distinct eigenvalue. As a result, it *must* be that $\underline{\underline{\Lambda}} \cdot \underline{A}$ is proportional to $\underline{A}$, since there is no other way that it could have the same eigenvalue. Let us call this constant of proportionality ω^2 (for reasons that should be clear). In this case, you should be able to see that we have:

$$\underline{\underline{S}} \cdot \underline{A} = \beta \underline{A} \qquad \Longrightarrow \qquad \underline{\underline{\Lambda}} \cdot \underline{A} = \omega^2 \underline{A}, \tag{20}$$

or in other words if $\underline{A}$ is an eigenvector of $\underline{\underline{S}}$, it is also an eigenvector of $\underline{\underline{\Lambda}}$!

2.2 Computing the normal modes and their frequencies

As can be seen from the previous section, the normal modes have already been computed for our problem, since they are just the eigenvectors of $\underline{\underline{S}}$! In order to compute their eigenvalues, all we need to do is compute

$$\underline{\underline{\Lambda}} \cdot \underline{A}^k = \omega_k^2 \underline{A}^k. \tag{21}$$

However, since we know that the first component of $\underline{A}^k$ is always 1 by construction, we just need to compute the first row of the matrix multiplication! Notice also that we have so far not needed to even begin constructing the $\underline{\underline{\Lambda}}$ matrix, and therefore our method is extremely general, and can be applied to any problem which has a discrete rotation symmetry described in the previous section.

Exercise: Show that the matrix $\underline{\underline{\Lambda}}$ for the 3-mass problem can be written as

$$\underline{\underline{\Lambda}} = -\omega_0^2 \begin{pmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{pmatrix} \tag{22}$$

Exercise: Use the above result to show that for the 3-mass problem, the eigenfrequencies satisfy the equations (with $N = 3$). Then, show that the equations below actually work for any N. (You will need to guess the form of the $\underline{\underline{\Lambda}}$ matrix in the case of N masses.

$$\omega_k^2 = \omega_0^2 \left(2 - \beta_k - \frac{1}{\beta_k}\right) = 2\omega_0^2 \left(1 - \cos\left(\frac{2\pi k}{N}\right)\right) = 4\omega_0^2 \sin^2\left(\frac{\pi k}{N}\right). \tag{23}$$

We have something quite interesting to note here: for all values of k except $k = 0$ and $k = N/2$ (which correspond to $\beta_k = \pm 1$) we have two different values of k which give the same normal mode frequency! For example, $k = 1$ gives the same normal mode frequency as $k = N - 1$, $k = 2$ gives the same frequency as $k = N - 2$ and so on! (This should make it clear why both $k = 0$ and $k = N/2$ – if it is a root – have been excluded.) Thus, while $\underline{\underline{S}}$ does not have any degeneracy, $\underline{\underline{\Lambda}}$ seems to have a lot of degeneracy! It turns out that this degeneracy is essential to solving the problem.

Exercise: Show that when $\beta_k = \pm 1$, i.e. when $k = 0, N/2$, the normal modes A^k have only real entries.

Exercise: Now show that for all $\beta_k \neq \pm 1$, i.e. all $k \neq 0, N/2$, A^k and A^{N-k} are complex conjugates. Both of these eigenvectors have the same eigenvalue ω_k^2.

Exercise: Show that if you have two eigenvectors $\underline{V}_1$ and $\underline{V}_2$ with the same eigenvalue ω^2, then any linear combination $a\underline{V}_1 + b\underline{V}_2$ also has the same eigenvalue ω^2.

Exercise: Now, the normal modes for a system of oscillating masses should be real numbers. The degeneracy allows us to do this quite simply. For every such pair, construct two *real* normal modes

$$\underline{B}^k = \frac{\underline{A}^k + \underline{A}^{N-k}}{2} \qquad \text{and} \qquad \underline{\tilde{B}}^k = \frac{\underline{A}^k - \underline{A}^{N-k}}{2i} \tag{24}$$

Exercise: Show that the normal modes of this system have three classes:

(a) A zero frequency normal mode (corresponding to $k = 0$) which looks like

$$\begin{pmatrix} 1 \\ 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix},$$

(b) The real parts of A^k, which look like

$$\begin{pmatrix} 1 \\ \cos 2\pi k/N \\ \cos 4\pi k/N \\ \cos 6\pi k/N \\ \vdots \\ \cos 2(N-1)\pi/N \end{pmatrix},$$

(c) The imaginary parts of A^k, which look like

$$\begin{pmatrix} 0 \\ \sin 2\pi k/N \\ \sin 4\pi k/N \\ \sin 6\pi k/N \\ \vdots \\ \sin 2(N-1)\pi/N \end{pmatrix}.$$

3 The infinite coupled system: discrete translation symmetry

We can now extend what we have done so far to an *infinite* system shown in Figure (4) which possesses a different – although related – symmetry: discrete *translation* symmetry. Consider an infinite system of coupled oscillators, coupled just as before to their nearest neighbours. Just as before, this system possesses a symmetry: if you went out of the room for a minute and came back in, and your friend shifted the entire infinite system by the distance between two masses (say, a) then when you returned you would have no way of knowing if such a shift had even been made. In other words, the system is symmetric under discrete translations by a distance a.

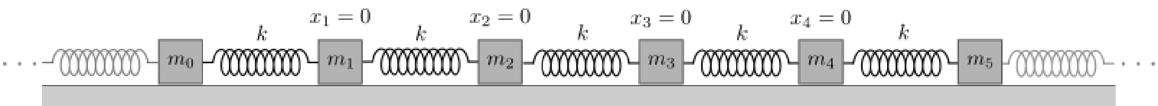

Figure 4: An infinite system of springs and masses: the masses are all identical $m_i = m$, and separated by a distance a. The masses are indexed by an integer j, and an arbitrary mass is indexed as $j = 0$.

As before, let us translate this into mathematics. We can describe our system – just as we did in the periodic case – using a vector comprising the positions of each of the masses. However, unlike the periodic case, this would now be an *infinite* vector. Thus, its entries would look something like

$$\underline{A} = \begin{pmatrix} \vdots \\ A_{-1} \\ A_0 \\ A_1 \\ A_2 \\ \vdots \\ A_N \\ A_{N+1} \\ \vdots \end{pmatrix},$$

where we have chosen one arbitrary mass to be the "zero-th" mass.[2] Now, what would our symmetry operation do to the system? Well, if we were to move the entire system leftwards by one unit a, then this would just be equivalent to moving $A_0 \to A_1$, $A_1 \to A_2$, and so on. In general, this would be equivalent to moving $A_j \to A_{j+1}$. The matrix that performs such an action is very similar to the $\underline{\underline{S}}$ matrix we saw in the previous section, and we will call it $\underline{\underline{T}}(a)$ to distinguish it from the previous $\underline{\underline{S}}$.

Exercise: Convince yourself that the form of the $\underline{\underline{T}}(a)$ matrix is

$$\left(\underline{\underline{T}}(a)\right)_{i,j} = \delta_{i,i+1}.$$

[2]Notice how the translation symmetry is basically a restatement of the fact that one could choose *any* mass to be the "zero-th" mass, since they are all identical.

Exercise: Now convince yourself that the form of the $\underline{\underline{K}}$ matrix is:

$$\begin{pmatrix} \ddots & \vdots & \vdots & \vdots & \vdots & \iddots \\ \cdots & 2k & -k & 0 & 0 & \cdots \\ \cdots & -k & 2k & -k & 0 & \cdots \\ \cdots & 0 & -k & 2k & -k & \cdots \\ \cdots & 0 & 0 & -k & 2k & \cdots \\ \iddots & \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \tag{25}$$

As before, since we have the symmetry matrix and know its action on any vector, we can use the eigenvalues of $\underline{\underline{T}}(a)$ matrix to find its eigenvectors, just as we did for the symmetry matrix in the previous section. Let us first guess the form of the eigenvectors. Remember that the action of $\underline{\underline{T}}(a)$ was to shift the entire system leftwards, and so – just as we had for the case of discrete rotations – we can say that the eigenvectors of $\underline{\underline{T}}(a)$ have components which satisfy

$$\underline{\underline{T}}(a) \cdot \underline{A} = \beta\, \underline{A} \qquad \Longrightarrow \qquad \underline{A}_{j+1} = \beta \underline{A}_j. \tag{26}$$

What can we say about β? Since there are an infinite number of masses, we do not possess the simplification that we did in the previous section, i.e. we cannot say that after some fixed number of iterations k, we'd get back the identity matrix. This is to say that for an infinite system, we can translate it an infinite number of times and still never get back to where we originally were, which should make intuitive sense. Therefore, we have no restriction on β, it can take *any* value, and has – associated with it – a unique eigenvector A^β (up to some overall multiplicative constant).

Exercise: Show indeed that $A_{j+1} = \beta A_j$, and $A_{j-1} = \beta^{-1} A_j$.

Thus we can show that the general form of A^β is:

$$\underline{A}^\beta = \begin{pmatrix} \vdots \\ \beta^{-1} \\ 1 \\ \beta \\ \beta^2 \\ \beta^3 \\ \vdots \\ \beta^N \\ \beta^{N+1} \\ \vdots \end{pmatrix}, \tag{27}$$

where we have chosen one of the components, say A_0, to be 1, and for all other j, $A_j = \beta^j$. Now, just as before, since the eigenvectors of $\underline{\underline{T}}(a)$ are also eigenvectors of $\underline{\underline{\Lambda}} = \underline{\underline{M}}^{-1} \cdot \underline{\underline{K}}$, we can find the eigen*values* of $\underline{\underline{\Lambda}}$, which – just as before – are the normal mode frequencies ω_β (notice how we add a subscript β to denote that each eigenvector possesses its own eigenfrequency).

Exercise: By multiplying row 0 of the $\underline{\underline{\Lambda}}$ with $\underline{A}^\beta$, show that

$$\omega_\beta^2 = \frac{k}{m}\left(2 - \beta - \frac{1}{\beta}\right). \tag{28}$$

We notice something interesting: since for (almost) every β, $\beta \neq \beta^{-1}$, we see that $\omega_\beta = \omega_{1/\beta}$. Thus, for each pair of numbers β and β^{-1}, there is a unique frequency ω_β^2.

Exercise: There are only two values of $\beta = 1, -1$ for which this is not true. Why do you think this is the case? Show that for these values, the frequency is

$$\omega_{\pm 1}^2 = 2k \mp 2k. \tag{29}$$

Exercise: $\beta = 1$ represents (again) the zero-frequency mode. What do you think it represents *physically*?
Answer: It represents a global translation of the system along the $x-$axis. Under such a translation, the energy of the system remains constant. It turns out that zero-frequency modes, symmetries of systems, and conservation laws are intimately linked to each other. You will learn about them much later in advanced courses in physics, although you might not be able to make the connection back to these "simple" mechanical examples.

Exercise: Now, go back and repeat the analysis for a system of coupled *pendula*, coupled by springs. Would such a system *also* possess a zero-frequency mode?

Exercise: Now, go back to the previous section and look at the zero frequency mode there. What do you think *that* mode represents?
Answer: There the zero frequency modes represent uniform rotations on the bead, at some constant speed, so that the centre of mass of the system remains conserved.[a]

[a]I am not completely sure about this yet. I will confirm this with someone more knowledgeable, in case there is some subtlety that I am missing!

The fact that every distinct β and $1/\beta$ have the same frequency will have profound implications when we move on to the next section and try to solve *finite* systems using the result for an infinite system. Note that this is a curious feature of only one-dimensional systems; higher-dimensional systems do not possess it, and are consequently much harder to solve in this method.

4 Solving a finite system using an infinite system

We now have a very interesting possibility: as we have seen, the infinite system offers us an abundance of modes: we find an infinite number of normal modes, parametrised by β, which could be *any* real number. Thus, our normal modes and their frequencies are uncountably infinite! Indeed it would appear that an infinite system possesses *every* possible frequency conceivable.[3] Now, let us see if we can use this embarrassment of riches to solve a more restricted problem, that if N masses with fixed walls at either end. Thus, we have N masses, numbered from 1 to N, with un-moving and rigid walls at either end.

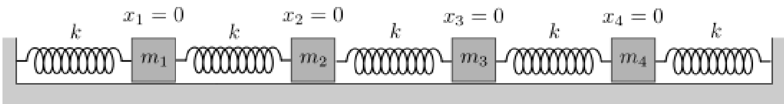

Figure 5: A finite system of springs and masses: again, the masses are all identical $m_i = m$, and separated by a distance a. The masses are indexed by an integer j, with j running from 1 to 4.

The idea is simple, but very deep:suppose for a minute that your system was actually infinite, but that you are only looking for the normal modes that kept the mass at 0 and the mass at $N+1$ always at rest. Remember that our interactions are *local*, meaning that every mass only knows what's happening to the masses on either side of it. This means that there is **no difference** between our finite system, and this *subset* of normal modes that keep the masses 0 and $N+1$ fixed.

Exercise: This is a complicated idea. Spend some time absorbing it.

So if you showed someone a video of N masses behaving a certain way, it would not be possible for them to distinguish between the motion of a finite system with fixed walls, or a combination of normal modes that keep masses 0 and $N+1$ fixed.

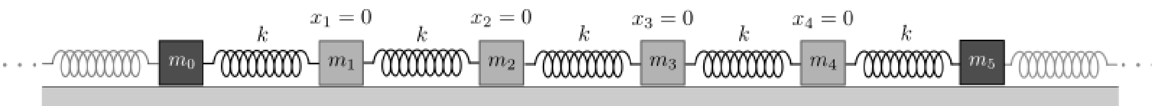

Figure 6: A finite system of springs and masses as a section of an infinite system: the masses at $j = 0$ and $j = N+1 = 5$ are held fixed artifically, and only those modes are considered which keep these two masses fixed.

Armed with this idea, we can now try to "filter" out those normal modes that keep the masses 0 and $N+1$ fixed in the infinite case. In order to do this, we need to define a new set of modes (call them $\underline{B}^\beta$) which have the same set of eigenfrequencies ω_β, but which also satisfy the conditions that $B_0^\beta = 0$, and $B_{N+1}^\beta = 0$. As we shall see, the reason that we can find such a new set of vectors is precisely because for every $\beta \neq \beta^{-1}$, there are two *different* eigenvectors with the *same* frequency ω_β.

Let us start by trying to satisfy the condition on the mass $j = 0$. Since we know $A_0^\beta = 1$, and $A_0^{1/\beta} = 1$. Thus, a good guess to get $B_0^\beta = 0$ is to define:

$$\underline{B}^\beta = \underline{A}^\beta - \underline{A}^{1/\beta}. \tag{30}$$

[3]If you are not a little disturbed by this, you have probably not understood it.

Exercise: Show that $B_0^\beta = 0$ for all β.

So far, we have not really made any restriction on β. Indeed, any value of β would give us $B_0^\beta = 0$. However, it is the *second* boundary condition that will give us a restriction on β.

Exercise: Show that

$$B_{N+1}^\beta = 0 \quad \Longrightarrow \quad \beta^{N+1} - \frac{1}{\beta^{N+1}} = 0.$$

Exercise: Using the above result, show that if we want the right-hand side to be zero, we require $|\beta| = |1/\beta|$. Of course, we are not discounting complex values of β, since we saw in the previous section that that was perfectly acceptable so long as we got real eigenvectors at the end. Use this to show that we can write $\beta = \exp(i\theta)$, where θ is some real number.

Exercise: Now, use the previous result to show that the general form of the vector $\underline{B}$ is:

$$\underline{B}^\theta = \begin{pmatrix} \vdots \\ 0 \\ e^{i\theta} - e^{-i\theta} \\ e^{2i\theta} - e^{-2i\theta} \\ \vdots \\ e^{i(N+1)\theta} - e^{-i(N+1)\theta} \\ \vdots \end{pmatrix} \propto \begin{pmatrix} \vdots \\ 0 \\ \sin\theta \\ \sin 2\theta \\ \vdots \\ \sin(N+1)\theta \\ \vdots \end{pmatrix}.$$

We see something very strange occurring: the different normal modes $\underline{B}$, now parametrised by θ instead of β, can be written as a column of different sines! Therefore, we can now ask ourselves the following question: what values can θ take, so that the $(N+1)-$th mass always stays fixed? Well, clearly, if

$$(N+1)\theta = n\pi \qquad \Longrightarrow \qquad \text{for} \quad \theta_n = \frac{n\pi}{N+1}, \tag{31}$$

this is satisfied. (It's also satisfied when $\theta = 0$, but that's pretty boring.) Thus, for all these values of θ, the normal modes are those that keep the masses at 0 and at $N+1$ always fixed. Now, we can make two important cosmetic simplifications. First, since we are only interested in the system between masses 1 and N, we can ignore all other terms in the eigenvector $\underline{B}$, and second, we can use n to denote the different normal modes, instead of θ_n. As a result, we can say that the normal modes of such a system can be written as

$$\underline{B}^n = \begin{pmatrix} \sin\left(\frac{n\pi}{N+1}\right) \\ \sin\left(\frac{2n\pi}{N+1}\right) \\ \sin\left(\frac{3n\pi}{N+1}\right) \\ \vdots \\ \sin\left(\frac{Nn\pi}{N+1}\right) \end{pmatrix}. \tag{32}$$

Thus, the general expression for normal modes of a finite number of masses between two walls is given by sine functions, as shown in Figure (7), with the $j-$component being given by

$$B_j^n = \sin\left(\frac{jn\pi}{N+1}\right), \qquad j = 1, 2, 3, \ldots, N. \tag{33}$$

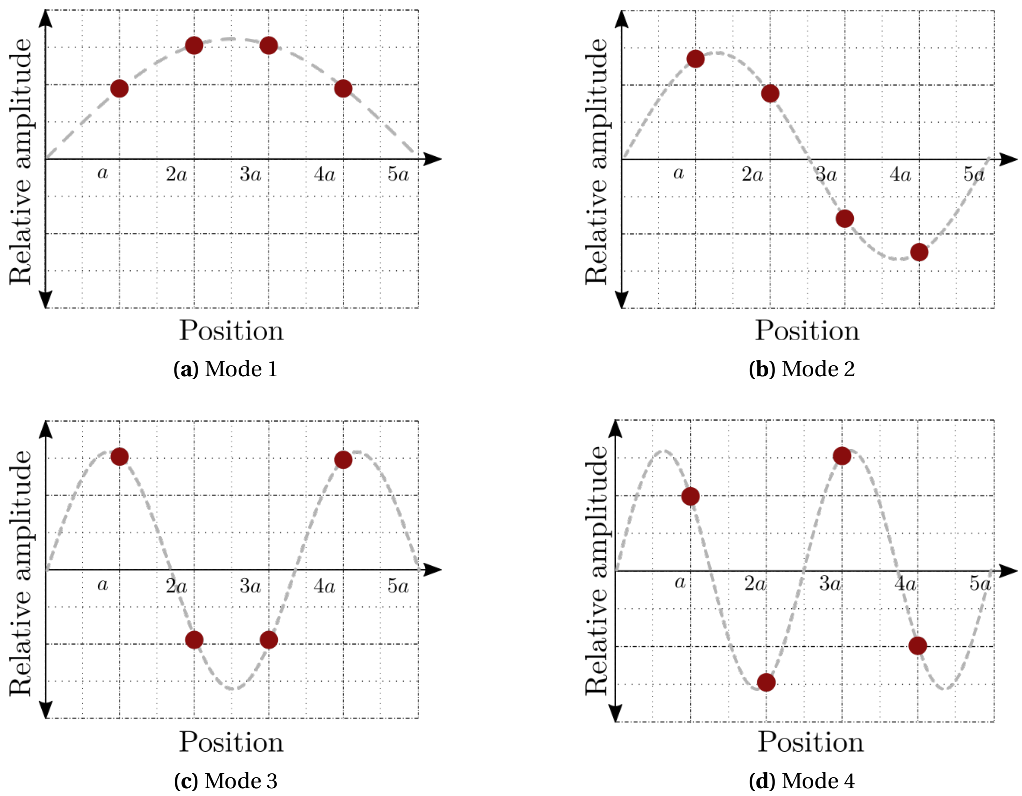

Figure 7: Relative amplitudes for the four different masses in each of the different normal modes. Each normal mode is associated with an "invisible" sine function with higher and higher frequency, with the amplitudes of each mass being given by the value of this sine function at the block's position.

We can now find the normal mode frequencies for these modes, using Equation (28).

Exercise: Show – using Equation (28) – that the normal mode frequencies can be written as:

$$\omega_n^2 = \frac{2k}{m}(1-\cos\theta_n) = \frac{4k}{m}\sin^2\left(\frac{\theta_n}{2}\right) \implies \omega_n = 2\sqrt{\frac{k}{m}}\sin\left(\frac{n\pi}{2(N+1)}\right). \tag{34}$$

Notice what the above equation is telling you: every normal mode, enumerated by n, has its own frequency ω_n (this much we know and expect, since that was the very definition of a normal mode), and apart from overall factors of k/m and so on, this frequency depends purely on the specific mode (n) and the total number of masses (N). Therefore, if we define $\omega_0 = \sqrt{k/m}$, we can say that $\omega_1 = 2\omega_0\sin(\pi/10)$, $\omega_2 = 2\omega_0\sin(2\pi/10)$, $\omega_3 = 2\omega_0\sin(3\pi/10)$, $\omega_4 = 2\omega_0\sin(4\pi/10)$: "higher" modes have higher frequencies.

Exercise: Define a "length" $L = (N+1)a$, and write

$$\omega_n = 2\sqrt{\frac{k}{m}}\sin\left(\frac{n\pi a}{2L}\right) = 2\omega_0\sin\left(\frac{n\pi a}{2L}\right). \tag{35}$$

You can now argue that the "wavelengths" of the different modes are given by $\lambda_n = 2L/n$. Using the definition of a wave-vector as $k_n = 2\pi/\lambda_n$, show that

$$\omega_n = 2\omega_0\sin\left(\frac{k_n a}{2}\right). \tag{36}$$

Such a relation between ω and k is known as a **dispersion relation**, and we will see such relations very often during the rest of this course.

Notice that the wavelength of the modes k_n and the frequencies of the modes ω_n are not linearly related, they are related by a *sine* function. What this means is that the periodicity in *space* (as shown in Figure (7)) and the periodicity in time (given by Equation (36)) are not linearly related to each other. As a result, the quantities

$$\frac{\omega_n}{k_n} \quad \text{and} \quad \frac{\mathrm{d}\omega_n}{\mathrm{d}k_n}$$

are not equal. This might not seem like an important thing, but you should come back and look at this problem when we begin to study waves on a continuous string.

4.1 Another example: beads on a string

The entire analysis we've done so far could just as well be extended to a slightly different – although mathematically identical – problem, that of beads on a string. The analysis for such a system is mathematically identical to the analysis we did for masses connected by springs.

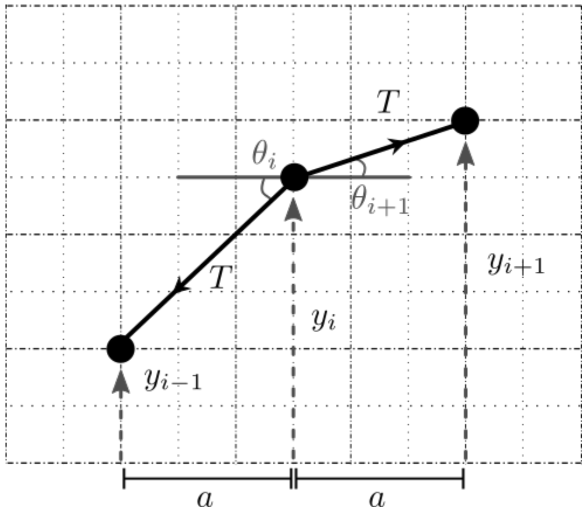

Figure 8: A beaded string: the masses, each with mass $m_i = m$ and with coordinates x_i and y_i, are spaced a distance a apart and connected by a massless string, held together by some tension T. We assume small oscillations, meaning that the angles $|y_i| \ll a$, and so $\theta_i \ll 1$.

The angles θ_i and θ_{i+1} are given by

$$\begin{aligned} \tan\theta_i &= \frac{y_i - y_{i-1}}{x_i - x_{i-1}}, \qquad \Longrightarrow \qquad \theta_i \approx \frac{y_i - y_{i-1}}{a}, \\ \tan\theta_{i+1} &= \frac{y_{i+1} - y_i}{x_{i+1} - x_i}, \qquad \Longrightarrow \qquad \theta_{i+1} \approx \frac{y_{i+1} - y_i}{a}, \end{aligned} \tag{37}$$

where we have used the small-angle approximation, i.e. that $\tan\theta \approx \sin\theta \approx \theta$, since the angles are small.

Exercise: Show by applying Newton's laws to the bead i that the equations of motion in the y and x directions are given by:

$$\begin{aligned} m\ddot{y}_i &= -T_i \sin\theta_i + T_{i+1}\sin\theta_{i+1}, \\ m\ddot{x}_i &= -T_i \cos\theta_i + T_{i+1}\cos\theta_{i+1}. \end{aligned} \tag{38}$$

Now, we assume that the motion is only in the transverse direction, meaning that the beads have no horizontal acceleration. This means that

$$T_i \cos\theta_i = T_{i+1}\cos\theta_{i+1}, \tag{39}$$

and if we further use the small angle approximation to say that $|\theta| \ll 1 \implies \cos\theta \approx 1$, we see that

$$T_i \approx T_{i+1} = T, \tag{40}$$

i.e. the tension is constant in the string.

Exercise: Now, using the equation for the transverse motion and the small-angle approximation, show that:

$$\ddot{y} \approx \frac{T}{m} \times (\theta_{i+1} - \theta_i) \tag{41}$$

Exercise: Next, write the angles in terms of the transverse coordinates y to show that

$$\ddot{y}_i = \omega_0^2 \left(y_{i-1} - 2y_i - y_{i+1} \right), \tag{42}$$

where $\omega_0^2 = T/ma$.

Clearly, the acceleration of the ith mass on the bead is identical in form to the acceleration of the ith mass in the spring-mass system! As a result, no more calculation is needed, and one can immediately quote the results of the previous section to describe the movement of the beaded string.

What's even better, the configurations given in Figure (7) would still hold true. However, while earlier the red dots represented relative amplitudes, in this case, we can actually use them to represent each bead. Thus, each dot would show the the *position* coordinate of each of the beads on the string. Therefore, the same Figure (7) would also represent a "snapshot" of the string in the different normal modes!

Exercise: Add straight lines – representing segments of the string – in between the "masses" shown in Figure (7) to get the configuration of the beaded string in the different normal modes.