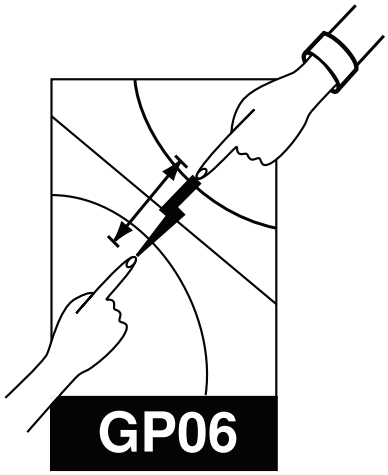
Harmonic motion and normal modes



edited PGH/PGR/AL/CP/JJL August 2022

1 Introduction

Simple harmonic motion (SHM) refers to the sinusoidal oscillations experienced by an object acted on by a restoring force directly proportional to its displacement. You are probably already familiar with the behaviour of single bodies – with a single mode of oscillation and a single natural frequency – undergoing SHM, such as a simple pendulum in a gravitational field. You might also be acquainted with the oscillations of extended, continuous bodies with an infinite spectrum of modes and frequencies, such as an idealised taut string oscillating in one of its harmonics. These two limits of (i) one and (ii) infinitely many modes of oscillation are explored in experiments GP04 — *Driven harmonic motion*[1] and GP03 — *Resonances and harmonics*[2], respectively, and form an integral part of the first-year physics syllabus.

You may be less familiar with the 'intermediate' regime, where a small number of discrete masses are coupled together by harmonic interactions and allowed to oscillate. Such systems possess a finite set of modes of oscillation called *normal modes*. The oscillations of a triatomic molecule provide a vivid illustration of this physics. As shown in Figure 1 below, this molecule can exhibit *twisting*, *rocking*, and *wagging* modes, depending on how it is excited. By studying systems in this intermediate regime with finitely many normal modes, you will better understand how to 'bridge the gap' between single oscillating bodies and continuous oscillating media.

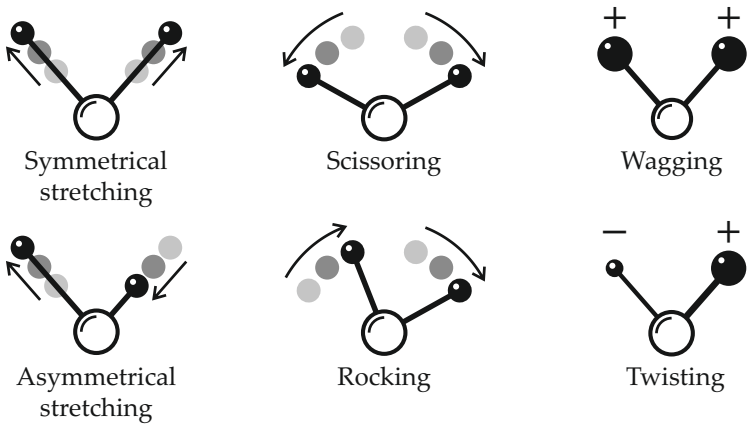


Figure 1: The six normal modes of a bent triatomic molecule (such as H_2O or ClO_2).

In this practical, you will study the normal modes of a one-dimensional system of coupled masses. Don't worry if you haven't met the relevant theory yet; the essential formulae are summarised below.

[1] www-teaching.physics.ox.ac.uk/practical_course/scripts/srv/local/rscripts/trunk/General/GP04/GP04.pdf
[2] www-teaching.physics.ox.ac.uk/practical_course/scripts/srv/local/rscripts/trunk/General/GP03/GP03.pdf

 Copyright © 2025 University of Oxford, except where indicated.

2 Theory

2.1 Oscillation of a single mass

In the first part of our experiment (Section 4.1), we will investigate the oscillations of a single mass connected by springs to two fixed points. The theory of the free oscillation in the direction x of a *single* body mass m under the influence of a restoring force $-kx$ is described in GP04 — *Driven harmonic motion*[3]. If, as is the case in Section 4.1 of this experiment, the mass is constrained by two springs, one can simply replace k by $2k$ everywhere. Thus the equation of motion becomes,

$$m\frac{d^2x}{dt^2} = -2kx. \tag{1}$$

As you probably know, the most general solution to this equation is

$$x(t) = A\cos(\omega t - \phi) \tag{2}$$

where $\omega = \sqrt{2k/m}$.

2.2 Oscillation of two equal masses

In the second part of our experiment (Section 4.2), we will consider two equal masses, shown in Figure 2.

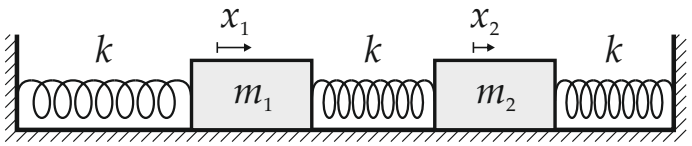


Figure 2: Schematic representation of the two-mass experiment.

Consideration of the effect of a connecting spring with an extension equal to $x_1 - x_2$ leads to a pair of simultaneous differential equations describing the acceleration of the two masses:

$$\begin{aligned} m_1\ddot{x}_1 &= -kx_1 - k(x_1 - x_2), \\ m_2\ddot{x}_2 &= -kx_2 - k(x_2 - x_1). \end{aligned} \tag{3}$$

where we ignore the small effect of damping for simplicity. If $m_1 = m_2$ then the equations are easily solved by a) adding and b) subtracting them to yield:

$$\begin{aligned} \frac{d^2}{dt^2}(x_1 + x_2) &= -\frac{k}{m}(x_1 + x_2), \\ \frac{d^2}{dt^2}(x_1 - x_2) &= -\frac{3k}{m}(x_1 - x_2). \end{aligned} \tag{4}$$

So, although the mass coordinates x_1 and x_2 do not *individually* obey the equations of motion for SHM, we can construct linear combinations of these coordinates that *do* oscillate harmonically. These linear

[3] www-teaching.physics.ox.ac.uk/practical_course/scripts/srv/local/rscripts/trunk/General/GP04/GP04.pdf

combinations [$(x_1 + x_2)$ and $(x_1 - x_2)$ in this case] are referred to as *normal mode coordinates*. The former coordinate oscillates relatively slowly with angular frequency $\omega_1 = \sqrt{k/m}$, while the latter oscillates more quickly at $\omega_2 = \sqrt{3k/m}$:

$$\begin{aligned}[x_1 + x_2](t) &= 2A_1 \cos(\omega_1 t - \phi_1),\\ [x_1 - x_2](t) &= 2A_2 \cos(\omega_2 t - \phi_2),\end{aligned} \tag{5}$$

where A_1, A_2, ϕ_1, and ϕ_2 are determined from initial conditions (the reason for arbitrary introduction the factor of 2 will become clear shortly). To calculate the behaviour of x_1 and x_2 individually, we simply find the sum and difference of the normal-mode coordinates to get

$$\begin{aligned}2x_1 &= 2A_1 \cos(\omega_1 t - \phi_1) + 2A_2 \cos(\omega_2 t - \phi_2),\\ 2x_2 &= 2A_2 \cos(\omega_1 t - \phi_1) + 2A_2 \cos(\omega_2 t - \phi_2).\end{aligned} \tag{6}$$

This can be elegantly rewritten as

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \underbrace{A_1 \begin{pmatrix} 1 \\ 1 \end{pmatrix} \cos(\omega_1 t - \phi_1)}_{1^{\text{st}}\text{ normal mode}} + \underbrace{A_2 \begin{pmatrix} 1 \\ -1 \end{pmatrix} \cos(\omega_2 t - \phi_2)}_{2^{\text{nd}}\text{ normal mode}}. \tag{7}$$

What does this mean? It tells us that, in general, the motion of each individual mass is a superposition of simple harmonic motion taking place at two different frequencies ω_1 and ω_2, the frequencies of the so-called *normal modes* of the system. The relative amplitudes of the normal modes (expressed by the constants A_1 and A_2) are set by the initial conditions. In fact, it is always possible to 'isolate' individual modes by choosing the initial conditions appropriately. For example, if you set up the masses at $t = 0$ at rest such that $x_1(0) = x_2(0) = A$, it necessarily follows that

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = A \begin{pmatrix} 1 \\ 1 \end{pmatrix} \cos(\omega_1 t). \tag{8}$$

Thus, each mass individual executes SHM at a fixed frequency ω_1 and with a fixed phase relationship relative to its neighbour [in this instance, the two masses are in phase ($x_1 = x_2$)]. The system is said to be 'in' its 1st normal mode. We can similarly isolate the 2nd normal mode by instead choosing $x_1(0) = -x_2(0)$; in this case, the masses move in perfect *anti*-phase ($x_1 = -x_2$).

More generally, a normal mode is a motion of the coupled masses where each the ratios of the displacements of the masses are fixed. For two masses, the configuration of the system in one of its modes can be written as

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = A \begin{pmatrix} a_1 \\ a_2 \end{pmatrix} \cos(\omega t) \tag{9}$$

where a_1 and a_2 are fixed coefficients for a given mode, Aa_1 is the amplitude of the motion for mass 1 and Aa_2 is the amplitude of the motion for mass 2. For the lower-frequency mode we have $a_1/a_2 = 1$ (fixed), while for the higher-frequency mode we have $a_1/a_2 = -1$ (fixed).

However, the motion of the system for arbitrary initial conditions is *not* SHM, and is given by a superposition (i.e., linear combination) of the two normal modes:

$$\begin{aligned} x_1 &= A_1\cos(\omega_1 t - \phi_1) + A_2\cos(\omega_2 t - \phi_2) \\ x_2 &= A_1\cos(\omega_1 t - \phi_1) - A_2\cos(\omega_2 t - \phi_2) \end{aligned} \tag{10}$$

In this practical, you will be able to measure the mode frequencies ω_i and the amplitudes A_i directly.

2.3 Adding more masses

In Section 5, you will examine the case of 3 and 5 gliders of equal mass, connected by identical springs.

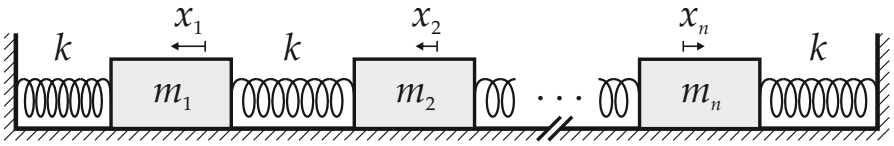


Figure 3: Schematic representation of the n-mass experiment.

Just as we had $n = 2$ normal modes for the 2-glider system, we should anticipate $n = 3$ and $n = 5$ modes in the 3- and 5-glider systems, respectively. One can show that the mode frequencies of a system of n coupled equal-mass gliders are given by

$$\boxed{\omega_j = 2\sin\left[\frac{j\pi}{2(n+1)}\right]\sqrt{\frac{k}{m}}\ (\text{for } j = 1\ldots n).} \tag{11}$$

The corresponding formula for the amplitude of the motion of the rth mass is,

$$\boxed{a_r = A_j\sin\left(\frac{rj\pi}{n+1}\right)} \tag{12}$$

where r labels the mass, j the mode. You will learn how to derive Eqs. (11,12) in your Wave Motion maths course. This practical will give you an opportunity to test both of these equations.

We can verify that these formulae reproduce the expected values for $n = 2$. For the frequencies, we get

$$\begin{aligned} \omega_1 &= 2\sin\left(\frac{\pi}{6}\right)\sqrt{\frac{k}{m}} = \sqrt{\frac{k}{m}}, \\ \omega_2 &= 2\sin\left(\frac{2\pi}{6}\right)\sqrt{\frac{k}{m}} = \sqrt{\frac{3k}{m}}, \end{aligned} \tag{13}$$

as expected. The amplitudes expected in the first normal mode ($j = 1$) are

$$\begin{aligned} a_1 &= A_1\sin\frac{\pi}{3} = A_1\frac{\sqrt{3}}{2}, \\ a_2 &= A_1\sin\frac{2\pi}{3} = A_1\frac{\sqrt{3}}{2}, \end{aligned} \tag{14}$$

i.e., $a_1 = a_2$, as required. Try verifying yourself that we recover the expected amplitudes for $j = 2$.

3 The apparatus



Remember that you are responsible for your own safety when working in the laboratory. Always talk to a demonstrator or technician if you are concerned that anything may be unsafe.

The apparatus consists of between one and five gliders that are free to move without friction in one dimension along an air-bearing track. The position of one or two of the gliders are measured by a motion sensor. This measures the transit time for an ultrasonic pulse to be reflected back by a 'flag' on the glider.

Make sure you can identify the following:

- the air track with supply (set pressure at '2', or high enough for the gliders to move freely)
- 5 gliders, 12 banana connectors, 6 springs, 2 black flags, a fixed endstop, a movable endstop, a plastic pulley wheel, a thread with a loop at each end, a hanger, and set of brass weights
- two motion sensor units (on their own stands) connected via an interface box (PASCO Universal Interface 850) to the computer; make sure green LED on the front of the interface box is lit
- an electromechanical driver and its variable frequency sine-wave generator (these are only used in the last optional part of the experiment).

Ensure that the apparatus is laid out along the bench in the following order: computer; air track supply; motion sensor (front) with its grille looking along the air track; the air track itself; motion sensor (back) at the rear of the track; the mechanical driver; and finally its power supply.

4 The experiment

4.1 Oscillations of one body

4.1.1 Setting up the track and measuring k of the spring

Switch on the air track supply. Move the free end of the track towards you so that it extends over the side of the bench (ensure the leg of the air track does not slip off the edge of the table!).

Put a glider on the track and adjust the levelling screws so that the glider does not accelerate. A bubble level is available to help ensure the track is truly horizontal.

Attach the fixed endstop at 0.00 m and the movable endstop at 0.80 m. Place a glider between them and attach it to the movable stop with a spring and two banana plugs.

Be careful not to push the banana plugs home into the holes.

Place the pulley wheel on the fixed endstop and run the thread over it from the glider to the weight hanger. Measure the extension of the spring for five or six different masses on the hanger. Plot your results to derive the spring constant k expressed in Newtons per metre.

Make spot checks to show that k for each of the other springs is the same within some tolerance.

Measure the mass of a glider on the laboratory scales.

Make the air-track parallel to the bench and re-level. Dismount hanger, thread and pulley, attach the

glider to the fixed endstop with a second spring.

4.1.2 Setting up the computer

Ensure that the sensor is plugged into the interface box. The yellow and black plugs should be connected to digital Channels 1 and 2, respectively. The interface box should be plugged in and connected to the computer via USB.

Log in to the computer. Open the 'PASCO Capstone' software. Click on the Hardware Setup button in the left-hand menu bar. You should see a picture of the interface box in the Experiment Setup window. On the picture, click the connection ports where the motion sensor is plugged in (Channels 1 & 2), and from the list of instruments choose Motion Sensor II. This sets the program up to read the motion sensor.

In the central part of the window, select Graph & Digits. The interface will change to display a large graph with two panels above it. Each of these has a Select Measurement button (for the graph this is on the axes). Select the panels to display position and time, and the graph to display position (it will automatically plot against time).

At the bottom of the window there are buttons to start/stop (and clear) the measurements and also to set the acquisition rate of the motion sensor — set the sample rate to 40 Hz.

4.1.3 Setting up the motion sensor

Clear all unnecessary objects from the bench along the length of the air track to avoid spurious reflection targets. Look along the track and align the motion sensor grille by observing your reflection in it. Pitch it up from the horizontal slightly to avoid reflections from the glider bodies themselves. Place one of the black reflection flags perpendicular to the track on the top of a glider and set it oscillating on a spring.

While the mass is oscillating, press the Record button in the bottom menu bar to take data in real time. You may need to scale the graph axes. This can be done by clicking on the axes digits and dragging.

Afterwards press Stop and the data is saved automatically.

You should adjust the angle of the motion sensor until a clean sine wave is shown on the display; it may be necessary to repeat a number of times to get the alignment correct. Ensure that clean signals are obtained to the end of the 2 m track.

4.1.4 Obtaining and analysing data

Start a new experiment, select File › New Experiment, and set up a motion sensor (as before, within Hardware Setup). Create an interface with two displays.

Place a 50 g weight on the glider and start it oscillating.

From the right-hand menu bar, drag the graph icon onto one display and drag the FFT icon onto the other. This will open a new graph displaying the FFT (fast Fourier Transform) of one of the Channels.

The data produced can be graphed and analysed in the PASCO Capstone software. In the icon bar above the graph, there is an option to fit curves to the measured data (select damped sine). The fitted curve will display on the graph, along with the curve fit parameters. Derive a value for the frequency and quality factor of the oscillation with errors. The quality factor, Q, determines how many oscillations the system makes before being damped to rest: the amplitude of the nth cycle falls off as $e^{-n\pi/Q}$.

To identify the dominant oscillation frequencies in your signal, you will use exploit its FFT. You will learn about the Fourier Transform of a continuous function in the second year; this is a related transform, the Discrete Fourier Transform (DFT). If you are interested, the maths is explained in detail in the

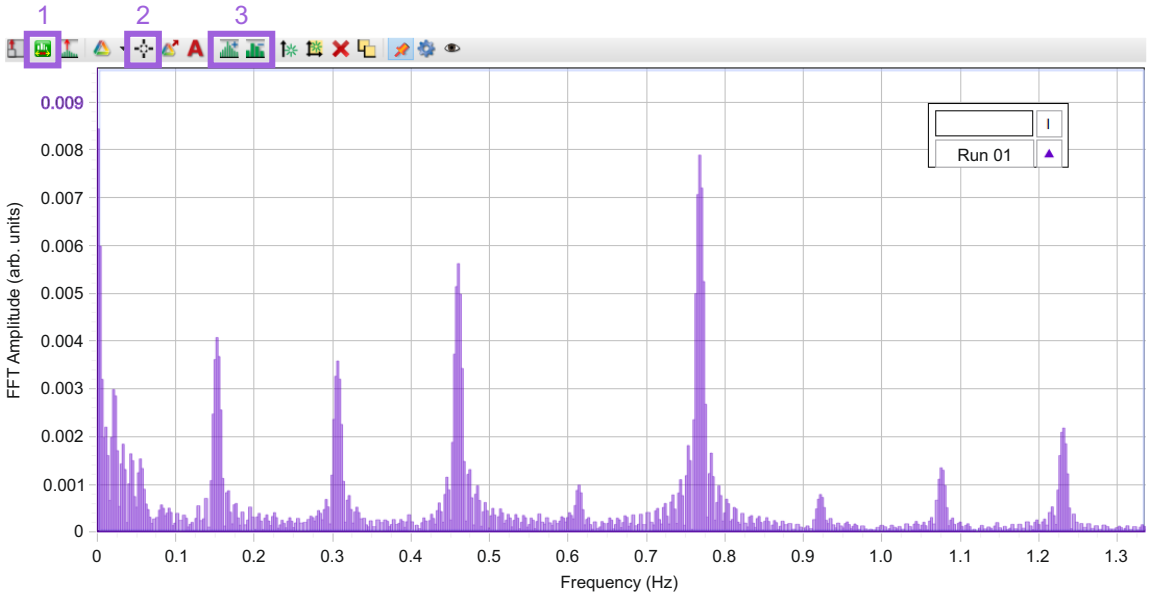


Figure 4: An illustrative FFT window in PASCO Capstone. Various tools are available above the graph: (1) Auto Sample Rate will automatically adjust the Sample Rate when you rescale the FFT window – it is advisable to **deactivate** this tool. (2) Coordinates Tool allows you to extract the coordinates of any data point. (3) Number of Bins allows you to more finely or coarsely sample the FFT.

Appendix. We have a set of N samples of position measurements $x(t_i)$ at equally spaced times $t_n = n\tau$, where, for the given sampling rate of 40 Hz, $\tau = 25$ ms. The DFT – which is essentially a set of oscillation amplitudes $a(f_j)$ – shows how we can construct the displacement $x(t_i)$ out of perfect harmonic functions (sines and cosines) with different frequencies f_j:

$$x(t_i) = \sum_j a(f_j) \exp(2\pi f_j t_i) \tag{15}$$

where the frequencies are also equally spaced according to $f_j = j/(N\tau)$. In order for x to come out real, we need both positive and negative f, with $a(-f_j) = a^*(f_j)$. Obviously we can't determine more than N values of a, so the frequencies must be limited to $|j| \leq N/2$; thus the maximum accessible frequency in the FFT for our sampling rate is 20 Hz. However, the relatively sedate motion of the gliders in this practical means that almost all of the signal of interest will be confined to frequencies below 1 Hz.

Look at the horizontal axis of the FFT graph you have created: you will see that it is a frequency axis, but with very high frequencies. You will want to rescale the axes, but before you do, **deactivate Auto Sample Rate** (see Figure 4 above) – this will prevent PASCO Capstone from automatically changing your sample rate from 40 Hz.

Once you've done this, stretch the frequency axis by dragging it until you have a maximum frequency of ~ 1 Hz. Click on the vertical axis and select the data from the active Sonic Ranger. Set the glider in motion and record x-data and a data for about 30 seconds. You should see a single peak at the frequency of the harmonic motion. You may need to set the binning to minimum bin size.

Print the page of data and FFT and read the frequency off the graph. Compare with the theoretical value of $(1/2\pi)\sqrt{2k/m}$.

▶ Compare the result for the frequency with theory based on your measured values of the spring constant and glider mass.

4.2 Coupled oscillations of two bodies

There are two sensors for this exercise, which can be used simultaneously to analyse the coupled behaviour of the oscillating masses.

To use the second sensor, make sure it is plugged into channels 3 and 4 on the interface box. Start a new measurement by selecting File ⟩⟩ New Experiment. As before, click on the Hardware Setup in the left-hand menu bar, but this time add both motion sensors (both pairs of cables to all 4 channels).

In the central part of the window, choose Two Displays. This will produce an interface with two graphs. Set up a position measurement for each motion sensor (click the button on each graph's y-axis and select Position — one each for channels 1 & 2 and 3 & 4).

The second sensor is mounted at the other end of the air track. Adjust its position so that it receives reflections from the second flag. The software will record data from both sensors automatically.

Double-check that the sampling rates for channels 1 & 2 and 3 & 4 are **both** 40 Hz at all times – if these rates differ, PASCO Capstone might fail to acquire data from the motion sensors.

Slide the movable endstop to 1.20 m, interpolating a second glider between it and the first glider connected with springs and banana plugs as before. Mount a sonic reflection flag and a 50 g weight on the second glider. Mount an extra sonic reflection flag on second glider. This extra flag may be attached with blue tack to the first and should protrude about 10 cm above it.

Adjust the vertical position of the motion sensor so that its grille is just above the height of the flag on the first glider. You should be able to choose which reflection to use by adjusting the pitch with respect to the horizontal.

Start the gliders oscillating. By adding the extra flag, the two gliders will not have the same mass. Compensate for this by adding small weights or blue tack to the first glider.

Set up the following cases in order and record normal mode frequencies seen in the FFT from two position sensors and two FFTs. Make sure you also record the frequency and amplitude of the peaks for later use.

(a) Two equal-mass gliders set oscillating by moving just one of the gliders and releasing.

(b) Two gliders as above but set oscillating by moving the gliders symmetrically in the same direction by equal amounts and releasing.

(c) Two gliders as above but set oscillating by moving the gliders equal amounts towards each other.

You should have learned something about the symmetry of the two normal modes: because the system is mirror symmetric about the centre, the two modes have definite mirror symmetry.

▶ Describe the symmetry of the low- and high-frequency mode: are the motions of the two masses mirror images of each other, or are the two displacements in the same direction? Is this consistent with the theory in Section 2?

5 More masses

Set up first three and then five gliders, all with equal mass, and record data with enough different initial conditions to see all the mode frequencies in each case. Experiment with different choices of glider and initial conditions for recording to determine the mirror symmetry of each mode. A good choice is to measure gliders 1 and 2 in the three-glider case, and gliders 1 and 3 in the five-glider case, followed by gliders 1 and 2. (Note glider 3 will remain stationary in mirror-symmetric modes).

To setup the PASCO Capstone software for this section:

(a) Create a new interface with four sections. Choose File » New Experiment. As before, click on the Hardware Setup and add both motion sensors (all 4 channels). Then select Four Displays in the main window.

(b) Arrange two displays to graph the position measurements for each motion sensor and set the other two displays to show the FFT for each motion sensor. (As before, drag icons from the right-hand menu bar onto the panels).

(c) Make sure **both** sample rates are set to 40 Hz.

(d) On the FFT toolbar (above each graph), click the gear icon, and set Auto to Off.

To set up an initial condition, mark the equilibrium positions of all the masses, then turn off the air and displace each mass by an appropriate amount from its equilibrium position. Turning on the air will initiate the motion.

You should try (at least) the following initial conditions:

- 'Random' displacements of all the masses. In this case, the motion should be a linear combination of all the normal modes.

  ▶ Are the ratios of displacements on different masses, as measured by the FFT method, consistent with Eq. (12)?

- Set the initial position of the masses so that the ratios of displacements are appropriate for a *single* normal mode, i.e., according to Eq. (12), so that this mode is isolated.

  ▶ Do masses oscillate in a SHM? Do you find a single peak in the FFT measurement?

When you are confident you have found all the frequencies of the normal modes:

▶ Plot all your mode frequencies for the $n = 1, 2, 3, 5$ cases against $\sin\left[\frac{j\pi}{2(n+1)}\right]$ [see Eq. (11)] and from the slope determine a best value for k/m. Does the fitted ratio agree with the expected value derived from earlier measurements of k and m?

▶ Write a brief summary of the experiment in your logbook. Discuss your results with a demonstrator as part of your assessment.

Optional

6 Optional section: Driven harmonic motion

If you have time, employ the mechanical oscillator to *drive* a system of two or three coupled masses.

In the 'PASCO Capstone' software, click on the Hardware Setup button in the left-hand menu bar. On the picture of the PASCO Interface box, click the connection port where the mechanical oscillator is connected (Output 2), and choose Output Voltage Sensor. Now click on the Signal Generator button in the left-hand menu bar. From here you can drive the signal generator using 850 Output 2.

You can now investigate the phase of a coupled mass compared to that of a driving oscillation for frequencies close to the resonance values established in your earlier work.

> Replace the end-stop with the mechanical oscillator (black cylinder mounted on a retort stand) and set up everything else as for the $n = 2$ or $n = 3$ case.
>
> Set up a graph to plot the position of the first mass and the Output Voltage of the driver versus time. You can show both on one plot by clicking on the y-axis and selecting Add y-axis. Ensure that you place the motion sensor such that positive displacement is in the same sense for both mass and driver (away from the mechanical oscillator).
>
> Start with all the masses at rest, and, in Signal Generator, set the oscillator to drive at the frequency of a normal mode that you have found previously. Save this plot.
>
> Now drive at a frequency just below the resonant frequency. Repeat this with a frequency just above the resonance frequency. Repeat for other resonance frequencies if you have time.

▶ How does the oscillation phase of the mass relate to that of the driver when the driving frequency is at, above and below the resonant frequency? Discuss any other interesting features you see.

A Discrete Fourier transform

A.1 Background information

The FFT (fast Fourier transform) function used in this experiment is a specific implementation of the discrete Fourier transform (DFT). The DFT is a systematic way of searching for periodic structures in a data set. Suppose we have a set of N data readings y_i taken at equal intervals δ of some other variable: in this case time, but it could also be a set of readings equally spaced in distance. Any periodic dependence on this other variable, with period Δ, will appear in the y_i as a repeating cycle with numerical period in i of Δ/δ. We extract this by setting up a set of cyclical data sets with different frequencies and expressing the actual data set in terms of them. This is most conveniently done with vectors: we define a data vector $\mathbf{y}$ as a column vector of the N values y_i, and a set of cyclical vectors with 0, 1, 2 ... $N/2$ cycles in them. For example, take $N = 8$. We can make a vector with zero cycles $\mathbf{c}_0$, and vectors with one cycle in

both cosine and sine versions $\mathbf{c}_1$ and $\mathbf{s}_1$:

$$\mathbf{c}_0 = \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} \quad \mathbf{c}_1 = \begin{pmatrix} 1 \\ 1/\sqrt{2} \\ 0 \\ -1/\sqrt{2} \\ -1 \\ -1/\sqrt{2} \\ 0 \\ 1/\sqrt{2} \end{pmatrix} \quad \mathbf{s}_1 = \begin{pmatrix} 0 \\ 1/\sqrt{2} \\ 1 \\ 1/\sqrt{2} \\ 0 \\ -1/\sqrt{2} \\ -1 \\ -1/\sqrt{2} \end{pmatrix}$$

Note that in both cases the cycle is one step short of repeating: that is the next in the sequence would be the same as the first.

We can then do the same with two and three cycles:

$$\mathbf{c}_2 = \begin{pmatrix} 1 \\ 0 \\ -1 \\ 0 \\ 1 \\ 0 \\ -1 \\ 0 \end{pmatrix} \quad \mathbf{s}_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \\ -1 \\ 0 \\ 1 \\ 0 \\ 1 \end{pmatrix} \quad \mathbf{c}_3 = \begin{pmatrix} 1 \\ -1/\sqrt{2} \\ 0 \\ 1/\sqrt{2} \\ -1 \\ -1/\sqrt{2} \\ 0 \\ -1/\sqrt{2} \end{pmatrix} \quad \mathbf{s}_3 = \begin{pmatrix} 0 \\ 1/\sqrt{2} \\ -1 \\ 1/\sqrt{2} \\ 0 \\ -1/\sqrt{2} \\ 1 \\ 1/\sqrt{2} \end{pmatrix}$$

When we come to four cycles, only the cosine version is possible: the sine version is evaluated at 0 each time:

$$\mathbf{c}_4 = \begin{pmatrix} 1 \\ -1 \\ 1 \\ -1 \\ 1 \\ -1 \\ 1 \\ -1 \end{pmatrix}$$

If we attempt to calculate $\mathbf{c}_5$ we find it is the same as $\mathbf{c}_3$, while $\mathbf{s}_5 = -\mathbf{s}_3$. However this is not very surprising as we already have a complete set of 8 vectors $\mathbf{c}_0$–$\mathbf{c}_4$ and $\mathbf{s}_1$–$\mathbf{s}_3$. Very conveniently we find that they are all orthogonal, and their norms are $N = 8$ in the case of $\mathbf{c}_0$ and $\mathbf{c}_4$, and $N/2 = 4$ for the others.

We can now write the data vector in terms of these basis vectors:

$$\mathbf{y} = \sum_{i=0}^{N/2} b_i \mathbf{c}_i + \sum_{i=1}^{N/2-1} d_i \mathbf{s}_i \tag{A.1}$$

The orthogonality makes the evaluation of the coefficients easy: we just take the inner product of the equation with one of the vectors, for example $\mathbf{c}_2^T \mathbf{y} = 4b_2$. Any periodicity in the data with period

$N\delta/n$ or frequency $n/(N\delta)$ will appear in either or both of b_n or d_n depending on its phase. Because of the redundancy whereby $\mathbf{c}_5 = \mathbf{c}_3$ etc, a higher frequency signal, above $1/2\delta$, will appear at a lower frequency. A signal not exactly at one of these frequencies will appear in all coefficients, but peaked around the relevant frequency range. Thus this transformation is a general tool for displaying any periodicity in the data.

The discrete Fourier transform uses not these vectors but N complex pairs of them $\mathbf{e}_j$:

$$\mathbf{e}_0 = \mathbf{c}_0 \quad \mathbf{e}_{\pm j} = \mathbf{c}_j \pm i\mathbf{s}_j \quad \mathbf{e}_{N/2} = \mathbf{c}_{N/2} \tag{A.2}$$

These are still orthogonal (in the usual sense for complex vectors $\mathbf{e}_j^\star \cdot \mathbf{e}_k = 0$) but also they all have the same normalisation $\mathbf{e}_j^\star \cdot \mathbf{e}_j = N$. The components are conveniently written in terms of complex exponentials: the k'th component of the j'th vector has value $\mathrm{e}^{2\pi ijk/N}$. Thus the projection onto these basis vectors reads

$$\mathbf{y} = \sum_{j=-(n/2)-1}^{n/2} b_j \mathbf{e}_j \quad \text{where} \quad b_j = \frac{1}{N}\mathbf{e}_j^\star \cdot \mathbf{y} \tag{A.3}$$

With this complex basis set the coefficients b_j are also complex, but with the relationship $b_{-j} = b_j^\star$. Note that this implies b_0 is real, but $b_{N/2}$ is also real, since this is the vector with alternating ± 1 like c_4 in the example above. This the DFT has N distinct parameters: b_0, $b_{N/2}$ and the real and imaginary parts of b_j for $j = 1$ to $(N/2) - 1$. The argument of the complex coefficients is related to the phase of the periodic structure within the data set, so for many purposes all the relevant information is in the modulus, which displays the total amount of structure in the data at that frequency, regardless of its phase. Thus the end result of a DFT is usually a plot of the values of the modulus of the coefficients b_j, which encodes periodic structure with frequency $j/(N\delta)$. (Note that since the periodic vectors only depend on N they can be pre-computed and assembled into a matrix to transform the vector $\mathbf{y}$ into the vector $\mathbf{b}$. This is essentially what the FFT does.)

A.2 Using the FFT

The FFT function continuously transforms the data in a specific 'first-in first out' buffer within the computer memory and displays the result on the graph on the screen. The buffer starts filled with zeros, and as data is taken they fill up the buffer. When it is full the buffer starts throwing out earlier data. Since the distance data oscillates around a fixed distance, perhaps 0.6 m, while the buffer is filling the most conspicuous feature of the data in the buffer is the step from 0.6 to zero, and this dominates the FFT, producing waves that run through the transform with a large peak at zero frequency. It's very conspicuous when this vanishes, so you should always run the data until this happens, even if the gliders have come to rest, as the replacement of zeros with 0.6 is still necessary. However once the buffer is full you *should* stop taking data, as the buffer is now simply replacing data from the start, with large amplitude oscillation, with current data with smaller amplitude.

You can change N by changing the frequency bin width: the width is $1/(N\delta)$ where δ is given by the digitisation frequency F. To put it another way, the bin width is F/N. The default appears to be $N = 512$, which at 40 Hz fills the buffer in about 13 s. You should change this to 1024, for which you need to run the data for about 30 s. The maximum is 2048, which requires a minute of data at 40 Hz, which is rather longer than the gliders will oscillate, so this is probably less useful.