0 Introduction

You may have heard of a 'torsional pendulum' in connection with the Cavendish experiment, in which very tiny forces of laboratory-scale gravitation can be detected. Torsional systems are not just curiosities, either, as many modern sensors and instruments make use of torsional deflection or torsional oscillation. But the TeachSpin Torsional Oscillator was designed to teach you not just about torsional systems, but also about a much more general class of motions and systems.

0.0 Harmonic Motion and its manifestations

The Torsional Oscillator is a mechanical system which will model for you an extremely broad class of other physical systems, all exhibiting the common feature of oscillations at a 'natural frequency'. You will perhaps first encounter the physics of this sort of 'simple harmonic motion' in the model system of one-dimensional linear motion of a point mass on a Hooke's-Law spring, but the physics you learn there will be applicable in much greater generality. TeachSpin's Torsional Oscillator achieves simple harmonic motion in a one-dimensional *angular* coordinate, rotation about a vertical axis.

The characteristic feature of harmonic oscillators is motion of sinusoidal form, for which the acceleration is not a constant, but is instead opposite to, and proportional to, the instantaneous value of the 'position coordinate'. In Hooke's-Law motion, this requires a non-constant force, which is a restoring force proportional to the displacement of the system. You'll find a similar force in the Torsional Oscillator, and you'll describe it in variables appropriate to rotational motion. But sinusoidal motion is so general that you'll want to recognize its features, and understand its complexities, even in non-mechanical systems.

Simple harmonic motion shows up in one-dimensional linear and rotational motion, but it also turns up in much more complicated mechanical systems, from bridges and buildings to crystals and molecules. Sinusoidal oscillations lie at the heart of wave motion in continuous systems, which is why they are fundamental to acoustics. Sinusoidal oscillations also appear in non-mechanical cases, such as the oscillations of charges and currents in electrical systems. You'll learn that even in empty space, electric and magnetic fields can undergo simple harmonic motion, describing the electromagnetic oscillations we call light. Beyond the borders of classical mechanics, you will find sinusoidal motion in the quantum-mechanical description of many systems. In fact, the harmonic variation of fields in general lies at the heart of quantum field theory, possibly our best hope for a 'theory of everything'.

So while you're about to encounter harmonic motion in the context of the Torsional Oscillator, what you will learn applies in many areas in physics, engineering, and beyond. It's safe to say that any system anywhere which displays oscillations at a characteristic frequency has been modeled as a harmonic oscillator.

0.1 Parts and names in the Torsional Oscillator

To familiarize yourself with your Torsional Oscillator, it's best to sit down in front of one, and handle it as you read through this introduction. The first thing you should know is that (unlike some delicate torsional oscillators built to detect tiny forces) the Torsional Oscillator before you is *robust*: you can touch and feel any piece inside it without fear of breaking anything.

You have a tall wooden ***case*** on a flat wooden ***base***, which ought to be level. Running the full height of the oscillator is a ***torsion fiber***, a length of strong steel wire. Find the upper segment of that wire, and pluck it like a guitar string -- you should hear the sound that results. Follow the fiber upward, and find the structure near the top front of the box that allows you to change the tension in the fiber. Adjust the tension, and confirm that you've made a difference by the 'plucking test'. Your oscillator is intended to work with enough tension in the fiber to yield a low musical note when the fiber is plucked.

Half-way down the fiber you can find the ***rotor*** structure, which is made of several pieces. Find the aluminum ***rotor shaft***, and see where its top and bottom are coupled to the fiber with ***wire clamps***. Near the top of the rotor shaft is a large ***rotor disc*** of pure copper -- go ahead and touch it, and set it into rotation with a twist. Its entire mass is being supported by the tension in the fiber. Near the bottom of the rotor shaft, you'll see another disc attached to it, and rotating with it. This is the rotor of the ***angular position transducer*** -- you can see it as the middle layer in a sandwich of three green fiberglass printed-circuit boards. Confirm that the middle layer does rotate relative to the stationary upper and lower boards fixed to the box. This sandwich, and the electronics attached to it, constitute the angular position transducer of your oscillator.

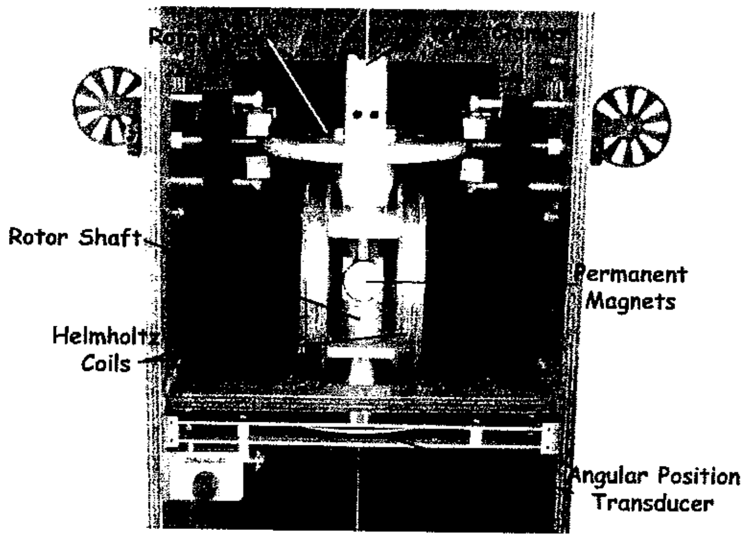

Fig. 0.1a: The rotor structure, with some of its parts labeled

At the center of the rotor shaft you can see some flat discs mounted to it -- these are some strong ***permanent magnets***, and they're in place to interact with the ***Helmholtz coil*** system wound on two black plastic bobbins. You'll use this magnet-in-coil system first as an angular velocity transducer, and later as a torque drive, for your oscillator. **Meanwhile, take care not to put magnetic materials near these strong magnets on the rotor.**

Turn your eye back to the copper rotor disc, and look for some curious 'disc brakes' that seem to be interacting with its periphery, left and right. These are the ***magnetic dampers***, which are mounted to the two sides of the wooden box. Find the knobs that are used to bring the dampers closer to, or farther from, the copper disc. Confirm that when the dampers are moved maximally inwards, they do not contact the copper disc, but give a bit of clearance above and below the copper. Find the brass thumbnuts on the outsides of the case that allow a vertical adjustment of the dampers, so as to achieve this clearance. For now, you might want to use the knobs to withdraw the dampers maximally outwards from the copper rotor disc. Note that the dampers also include some strong permanent magnets, **so keep magnetic materials away from their jaws.** (But copper is not magnetic -- no problem there.)

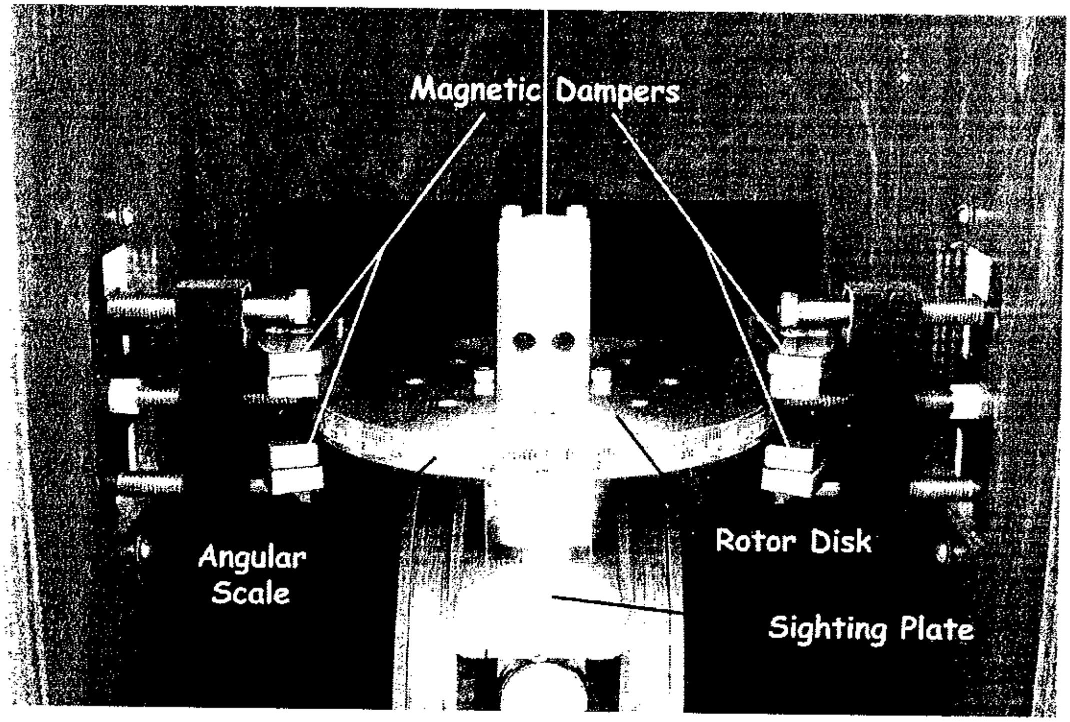

Fig. 0.1b: The top end of the rotor, with some parts labeled

Finally, on the periphery of the copper rotor disc, find the ***angular scale***, labeled with RADIANS at both ends, and with major markings numbered 1.0 through 5.0. In front of this scale, find the ***sighting plate***, with two fine lines drawn at its top center. Still the motion of the rotor with your fingers, and sight through the sighting plate, eyeballing through an overlapped view of the two sighting lines to a reading of the angular scale. You'll use this scale to calibrate the rotational motion of the rotor. Do not worry, at this stage, if the rotor settles down to give an oddball number like 2.98 on this angular-position scale.

0.2 Checking out operation of the Oscillator

This section assumes you have a Torsional Oscillator on a table in front of you, and that you're familiar with its parts (see 0.1). The goal of this section is to assure you that your Oscillator is working right.

Confirm that the wooden base is attached to the case, and that an AC line-power cord is attached to the back wall of the case. Find the power switch in the power entry module. Plug in to available AC line power, and turn on the switch. Now in the top left of the panel at the front bottom of the instrument, you should see a green LED light up to tell you power is on. This also energizes the electronics of the angular position transducer.

You can still touch any part of the Oscillator, now that it's on. Go ahead and try setting up a rotational motion of the rotor, by handling the copper rotor disc. If the magnetic dampers are fully withdrawn, you should see an oscillation that continues for well over a minute. If you hear any scraping sounds, or if the oscillation dies away in seconds, there is excess damping somewhere -- see section 0.3 for tips on aligning the oscillator. It's designed to allow rotational motion with no sliding contact anywhere -- look around at all the places where the rotor structure is near, but not touching, the stationary parts of the system.

Still the rotational motion by hand, and attach a voltmeter to the BNC output on the panel that's labeled ANGULAR POSITION. You should see some reading, which will only be stable if the rotor is truly at rest. Find the little brass panel under the green circuit boards with the ZERO ADJUST indication, and use this ten-turn knob to bring the voltage reading to zero. This is a convenience, to put zero signal at zero departure from equilibrium. Now grasp the edge of the copper disc, and give it a static deflection of about +1 radian -- use the angular position scale on the disc's periphery to show this. You should observe a change in the voltmeter reading, of order 1.5 - 2.0 DC Volts. Change the angular deflection to -1 radian, i.e., go to the other side of equilibrium. You should see a similar voltage with opposite sign. This is a first test that the angular position transducer is working -- section 1.2 will show you how to calibrate it.

If you're familiar with an oscilloscope, connect it to the angular-position output, and confirm that static deflections, and free oscillations, of the rotor will give the voltage signals you expect. (Choose **'DC coupling'** on your 'scope inputs, for viewing these low-frequency signals.) Some lovely oscillatory graphs should emerge. If you have a 2-channel 'scope, connect its other input to either (of the two) angular-velocity outputs on the Oscillator's front panel. Arrange for two-channel display on the 'scope, and see two sinusoidal waveforms. The angular-velocity output should also be generating a small but sinusoidal voltage waveform. Temporarily turn off the power to the Torsional Oscillator, hand-excite its rotational motion, and confirm that (though the angular-position signal disappears) the angular-velocity transducer still works. (How can this be?)

Now find a wire-clamp that's clamped to the torsion fiber **alone**, located in the lower chamber of the instrument's case. That's in place to give you a better way to excite the rotational motion of your Oscillator than by handling the copper disc directly. Use fingertip-and-thumb contact with the ends of this clamps, and an eyeball view of the copper disc, to pump energy into the rotational motion of the system. If you've ever ridden a playground swing, you'll know instinctively how to 'pump up' the system. In fact, tactile feedback from the wire clamp to your fingertips would let you achieve this even with your eyes closed. You've discovered the concept of resonance when you've found that you have to apply the fingertip pressure at the right frequency (the *oscillator's* choice of frequency) to achieve this pumping. Continue pumping until you have an oscillation that departs by ±1.5 or even ±2 radians from equilibrium -- **but don't go farther, lest you damage the torsion fiber.**

If you have set up a large oscillation, now use the knobs to bring the magnetic dampers inwards, and view the damping they cause in the oscillatory motion. Bring them fully inwards, and you'll see really dramatic damping. When the system is fully damped, you'll hardly be able to 'pump it up' as before -- instead, go back to handling the copper disc, give it a radian of deflection, and let it go, and view the results.

Now you're familiar with the use and handling of the apparatus, and if things are working right, you can go on to section 1 to learn about calibrations and measurements that you can perform with it. If things aren't working right, see section 0.3 on alignment.

0.3 Aligning a Torsional Oscillator

This section will tell you more about the adjustments that you can make to the Torsional Oscillator, in case you are not getting free oscillations of the rotor. You'll want to be familiar with the parts of the Oscillator (see 0.1) and you should have tried out its operation (0.2).

For a good alignment of the Oscillator, it's probably best to remove the magnetic dampers entirely at first. Use the brass thumbscrews on the outside of the box to do this, and when you've removed them, you can withdraw each whole damper bodily. Set it down where the strong magnetic fields between its jaws can do no harm.

Now find the ***rotor set-up tool***, a piece of black fiberglass sheet that comes with your oscillator. It may already be in place, lying atop the back of the upper support shelf of the Helmholtz coils. Slide it forward on that shelf until it slots under the rotor shaft, and its two 'arms' emerge into view behind the sighting plate. Now use the tension adjuster at the top center of the box to slack the torsion fiber, and you should see the whole rotor structure settle down onto this set-up tool. Re-tension the fiber until the rotor is just about, but not quite, lifting off this set-up tool.

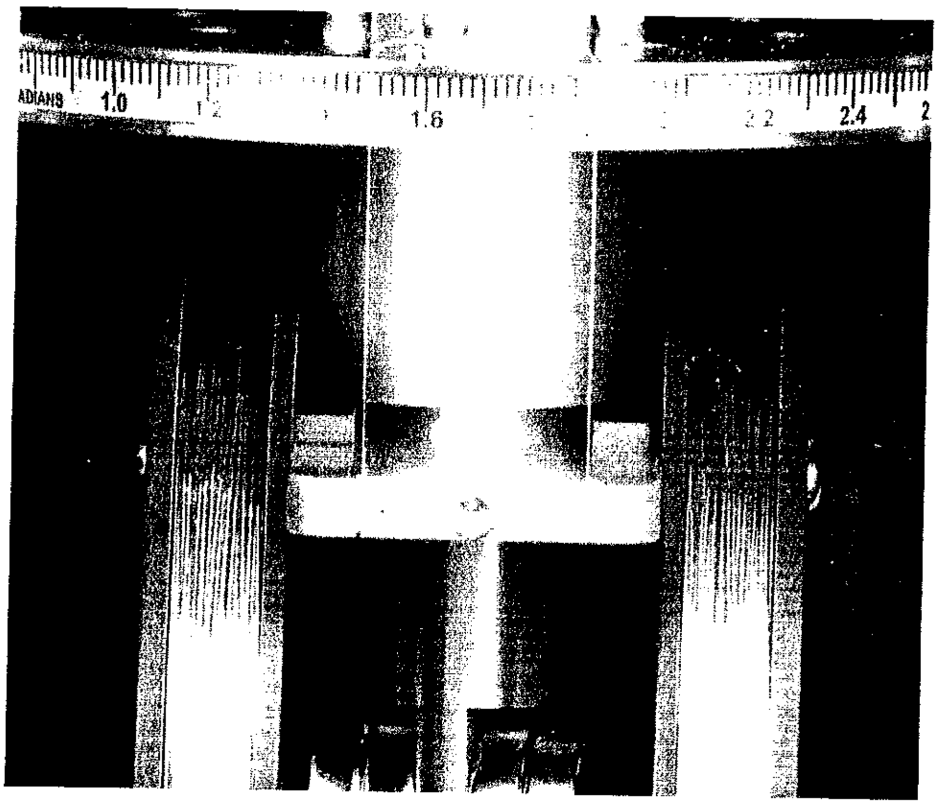

Fig. 0.3: The rotor set-up tool, in place supporting the rotor

You should now re-tension the torsion fiber until you see the rotor lift off the set-up tool; once it's free, slide the tool backwards on its support shelf to the back of the box. Continue tensioning using the 'plucking test', until you achieve a low musical pitch when you pluck the upper section of the fiber. The whole rotor structure should now undergo free torsional oscillations. Check that the copper disc is rotating in a horizontal plane, by comparing its vertical position behind the sighting plate. If you see the copper's plane wobble up and down by more than a mm, try adjusting the tightness of the four socket-head screws that attach the copper disc to the rotor shaft, so as to minimize this wobble. If you can't achieve rotation of the copper in its plane, see section 0.4 on (re)installing the rotor shaft on the torsion fiber.

Replace the magnetic dampers you removed, and use their knobs to withdraw fully the movable jaws of the dampers (watch the compression springs on two fixed shafts get compressed in this operation). When you mount the damping structures back into their places on the sides of the wooden box, see to it that two stainless dowel pins are pointing upwards (not downwards). Notice that each damper has, at the inner tip of the screw which rotates with its knob, a black nylon flat-headed screw. Adjust the vertical position of the whole damper structure, using the thumbnut adjustments, until this black screw head is vertically aligned with the copper disc of the rotor. Tighten down the thumbnuts.

Now check the horizontal clearances between the two black plastic screw heads and the angular position scale on the edge of the rotor. If they are unequal, here's how to center the rotor in the gap between them. Slack the fiber tension slightly. Notice the black bar that holds, at its center, the top end of the torsion fiber, and find the hinge, at the back of the box, about which this bar pivots. You will be able to slide the back end of this bar sideways, left or right, by several mm either way. Use this freedom of movement to slide the back of the bar, and hence the center of the bar, and hence the top of the fiber, and hence the whole rotor structure, laterally. After a trial adjustment, tension up the fiber and see it you've improved the lateral centering of the rotor disc. Iterate until you have equal clearance on both sides of the rotor.

Now that you've centered the rotor horizontally, turn your view to the position transducer near the bottom of the rotor shaft. Confirm that the middle layer of a 3-decker sandwich is lying parallel to, and halfway between, the upper and lower (stationary) decks. If it is not, you can adjust the position of those upper and lower decks. They're attached together, and attached to the wooden box by two aluminum 'card guides' at their left and right edges. If you need to, you can loosen the thumbnuts on the outside of the box that hold these card guides in place. Once they're loosened, you can adjust the location of the card guides vertically. Aim to adjust them, left and right, front and back, until you achieve a 3-layer sandwich of three parallel planes.

Having adjusted the vertical clearances of these transducer boards, you still can move them left and right, and forward and backward, it order to center them on the rotor. To accomplish the adjustments, you may need to loosen slightly the two 6-32 headless set-screws that snug them in place to the two aluminum bars holding them, at their sides, to the wooden case. Now you can slide the sandwich of upper and lower boards forward or backward in the slots in those two aluminum bars, as needed. To achieve the left/right adjustment, use an Allen driver to withdraw both set-screws a bit, and then drive inward with one of the screws until in pushes the whole board assembly sideways. when you've achieved the desired position, drive the other set-screw in gently until the board assembly is held between the two set-screws' pressure.

Finally, use the magnetic dampers' knobs to run the dampers' jaws fully inward, and adjust at the thumbscrew mountings until the dampers are aligned vertically to give equal clearances of the damping magnets above and below the copper disc. You should be able to get clearances of more than a mm between the copper and all eight of the rectangular permanent magnets in these dampers, and the clearance should persist even when the rotor is in a rotated position.

You may find, on occasion, that these damping structures stick in place when you use their screws so as to move them inwards. If they do, you can unstick them by hand intervention, and you can slide them out- and in-wards (against the springs' pressure) to exercise the lubricant on the two fixed shafts on which they ride.

0.4 Changing the torsion fiber

This section will teach you how to interchange torsion fibers in your oscillator. This will take a few minutes once you're experienced at it, but longer the first time. Since you might get a slightly different calibration constant for your angular-position transducer after this sort of interchange, you shouldn't do this swap unnecessarily. You'll need two Allen-wrench tools for these tasks -- each is called a 7/64" Allen driver, and we've supplied one L-shaped, and one with screwdriver handle, with the apparatus. These are the right size for turning all the 6-32 socket-head cap screws used in the wire clamps.

Changing a fiber is easiest if you can support the entire apparatus over a gap between two tables (see Figure 0.4). Give yourself a gap of 10 cm or so, and you'll have easy access to the bottom of the fiber, underneath the base.

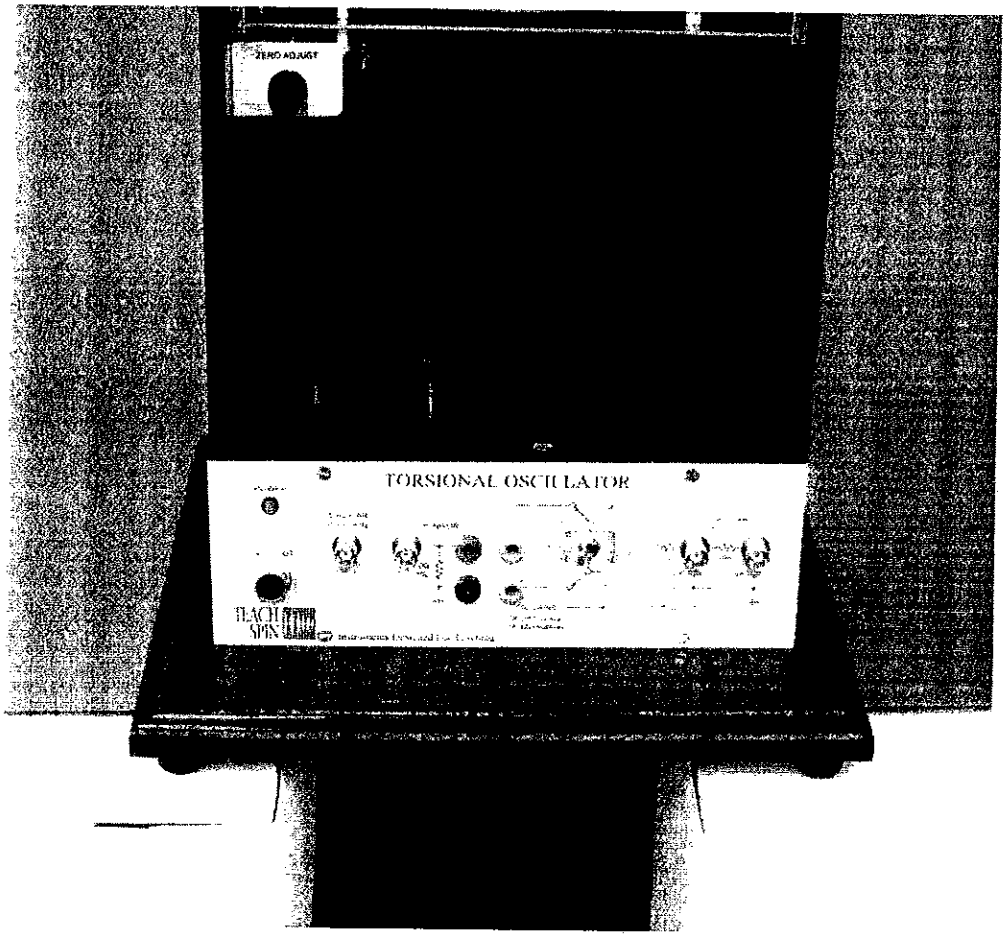

Fig. 0.4: The apparatus supported on two tables (for working on fibers)

Removing a fiber:

1. Put the rotor set-up tool (see section 0.3) in place to support the rotor structure from below, and now slack the fiber until the rotor just settles onto the set-up tool.

2. Loosen the two screws (with axes vertical) on the wire clamp that attach the clamp to the top of the rotor shaft. Next, loosen the two screws (axes horizontal) that hold the two halves of the clamp together. Now you can remove the two vertical screws, lift the clamp vertically, disengage it from the fiber, and set it aside.

3. Use a similar procedure to remove the clamp at the bottom end of the rotor shaft. This is harder to see, but it's identical in character to the clamp you just removed.

4. There may be another wire clamp attached to the torsion fiber (but to nothing else) on the lower section of the fiber. If it's there, take it off -- it's best to use both Allen drivers at once, so you can apply opposing torques to the two screws simultaneously, and avoid wrenching a twist into the fiber.

5. Now the torsion fiber is being held only at its top and bottom ends. Find the wire clamp at its top end, and remove it from the fiber as in step 2.

6. Crane your neck a bit, and get a view of the clamp that holds the bottom end of the fiber in place -- it's visible through a hole in the wooden base of the instrument. Leaving the clamp *attached* to the fiber, remove the two screws (axes vertical) that attach the clamp to the frame of Oscillator.

7. Pull the whole fiber, with bottom clamp still attached, vertically downward out of the instrument. This will leave the whole rotor structure, copper disc and all, wholly supported by the rotor set-up tool, which is in turn supported by the upper support plate of the Helmholtz coils. Leave the rotor there, since it's ready to receive the new fiber.

8. Get a close look at how the last remaining clamp is holding itself to the (bottom end of the) fiber. Go ahead and remove the clamp from the fiber.

Storing fibers:

At the back side of the wooden box, there's a clear plastic tube in place running along the centerline of the box. There should be three other fibers stored in that tube. You can slide off the red plastic cap at the top of the tube, and remove the whole tube from its holders. Swap the fiber you've removed for the fiber you want. See section 6.0 for advice on caring for these steel fibers, as they could rust if they get, and stay, damp.

Installing a fiber:

0. Have a look at one of the wire clamps you've removed. Note that it has two identical pieces, and learn how they fit together. Find the little V-groove in the working surface of the clamp half, and note that a used clamp may have that deformed a bit into a U-groove. The groove is there so

that the fiber will end up centered in the assembled clamp. It is best to clamp a fiber, **not between two V-grooves**, but between a V-groove and an originally-flat surface. You may reverse both halves of the clamps to engage a new fiber with heretofore-unused surfaces. You might want to use one set of surfaces for the two thicker, and the other set of surfaces for the two thinner, fibers.

1. Find the bottom clamp you removed, and attach it to one end of your new fiber. The fiber's end should be flush with the surface of the clamp, i.e., you want the fiber to be grasped by the clamp's full thickness. You do **not** need to tighten down the clamping screws until the two halves of the clamp meet! Instead, tighten down the two screws, symmetrically, to keep the clamping surfaces parallel, until you feel the steel begin to deform the aluminum clamping surfaces. Using the Allen screwdriver will limit the amount of torque you can exert to about the right level.

2. With the bottom clamp attached to the bottom end of the new fiber, it's time to thread this fiber upwards through the apparatus. **Unplug** the Oscillator from the AC power line before you do this! Peer underneath the wooden base to see that you get the fiber going vertically upward through the correct (largest) hole in the bottom anchor bar of the frame, and then see that it emerges through the cm-sized hole in the bottom wooden shelf of the Oscillator. Continue raising the top end of the new fiber until it goes upwards into the cm-sized hole in the bottom of the rotor shaft.

3. When you get the fiber's top end near the magnets on the rotor shaft, they'll conveniently grab and hold it magnetically. Now you have to get the fiber to go through a small hole, hidden from view, in the center of the rotor shaft. This is easiest to do by feel. Hold the fiber between finger and thumb, and roll your fingertips to rotate the fiber about a vertical axis, all the while exerting a modest upward lift on the whole fiber. Continue with this twirling motion until you feel the fiber mover upward vertically -- you're now 'through the narrows'. The magnet will continue to hold the fiber in place, but you can continue to slide it upwards.

4. As the fiber's top end nears the top anchor bar at the top of the box, guide it through the central hole in the top clamp, until it emerges. Finally now use two screws to attach the *bottom*-end wire clamp of the fiber to the bottom anchor bar -- more neck-craning to see how to do this.

5. The top end of the fiber is still loose. Before attaching its clamp, use the top tension-adjust knob to open up a gap of 6-10 mm of thread between the top anchor bar and the top hinge bar. This is to give you some room for the tensioning you'll do soon. Now attach the top wire clamp to the top end of the fiber, pulling as much fiber upwards through the clamp as you can before tightening the two halves together. Ensure that the fiber is caught in the V-groove on one side of the clamp. When you have the clamp attached to the fiber, attach the clamp to the black disc that's atop the top hinge bar. Ideally, you won't have to twist the fiber to make the screws line up with the holes.

6. The fiber is now anchored at bottom and top ends (though still free of the rotor structure). Now put a bit of tension into the fiber, to ensure it'll straighten out. Pluck the fiber to hear that you've achieved tension, and be sure you get a steady pitch for each new pluck. If the fiber is slipping in its end clamps, you'll hear descending pitches, and you'll have learned you weren't clamping the fiber tightly enough it its clamps.

7. With a modest tension in the fiber (less than you'll end up using), it's time to attach the whole rotor structure to the barely-taut fiber. Ensure that the rotor structure is still resting on the rotor set-

up tool, and ensure that it's at the angular orientation you'd like at equilibrium. The circular face of the permanent magnets should face right out at you, perpendicular to the axis of the Helmholtz coils. You needn't obsess on getting this perfect, but try to get within a few degrees of the right orientation. (It is *not* expected that this will put the 3.00 radian mark right behind the lines in the sighting plate.)

8. Mount a clamp to the fiber above the rotor structure, first well above the rotor, ensuring that the fiber is caught in the V-groove on one side of the clamp. When you have the clamp finger-tight on the fiber, slide it down the fiber until it seats on the rotor shaft's top end. Engage the clamp to the rotor with two screws, just finger-tight for now. Finish tightening the clamp to the fiber, and then the clamp to the rotor.

9. Repeat that procedure for the clamp at the lower end of the rotor shaft. Once you have the top and bottom clamps holding the rotor to the fiber, you are ready to finish tensioning the fiber. This should lift the rotor off of the rotor set-up tool, so that the rotor will oscillate freely. Slide the rotor set-up tool on its shelf to the back of the box when you're done.

10. See if you're content with the equilibrium orientation of the rotor that your set-up has given you. If you're off by more than 5 degrees (or 0.1 radians), reverse and repeat steps 7-9 until you get closer. If you're off by less than this, you can use the wire angular adjuster (at the top end of the fiber) to rotate the top, anchored, end of the fiber by $\pm 10°$, which will have the effect of rotating the equilibrium position of the rotor by about half this amount. (To use that angular adjuster, you may want to reduce the tension in the fiber. You'll need a 9/64" Allen wrench to loosen the two black-headed 8-32 screws that hold the knurled black angular-adjuster plate fixed. Rotate the plate about a vertical axis until you get the rotor to settle in the orientation you want, and then re-tension the fiber, and finally snug down those two black-headed screws.)

That's a long process in words, and it'll take a while your first time, but it soon gets quicker with experience. Once you have a new fiber in place, you should check section 0.3. on alignment, which will take only a minute once you're experienced. Recall that the rotating part of the angular-position transducer is likely to have ended up displaced by a mm or so in some direction in this whole process, so it'll certainly need re-zeroing, and might also have a slightly different calibration constant in your new assembly.

1 Operation and Calibrations

This chapter assumes you've worked through section 0.2, and have a working Torsional Oscillator before you. It describes many ways that you can experiment with the Oscillator, in order to learn the values of various parameters that can be used to model it.

1.0 Applying static torque

This section describes how to apply static torque to the rotor of the Oscillator using weights, and what you can do with this capability.

The rotor of your torsional oscillator responds to torque by accelerating angularly, just as in one dimension a mass responds to force by accelerating. If you apply a torque by direct hand contact to the copper rotor of your Oscillator, you can start it accelerating, but it'll reach some equilibrium angular deflection when the twist in the torsion fiber 'torques back' on the rotor to reduce the net torque (hand and fiber together) to zero. Here's a way to apply a torque to the rotor in a more quantitative way.

On the lower shelf of your Oscillator are stored two complete 'hang-down' units, each with a 50-g hanger and masses of 50, 100, and 200 grams (nominal). Also on the lower shelf are two low-friction pulleys, which can be mounted on vertical dowel pins found on the frame of your magnetic dampers. The last item you'll need is a piece (or two) of high-strength fishing line, useful for supporting the masses. Recall that a mass, *m*, supported at rest (or at constant velocity) in a gravitational field of strength, *g*, will require a support force of size, *mg*, and that this support force provided by a string will put the string under a tension of the same size.

Now you can arrange the line supporting the mass to be wrapped around the hub at the top end of your Oscillator's rotor shaft, and this automatically ensures that the line's tension acts perpendicularly to a 'lever arm', *R*, where *R* is the radius of the hub. There's a tiny screw-head at the top of the hub onto which the end of the string can be anchored. And that combination yields a torque of size *Rmg*, which will twist the rotor.

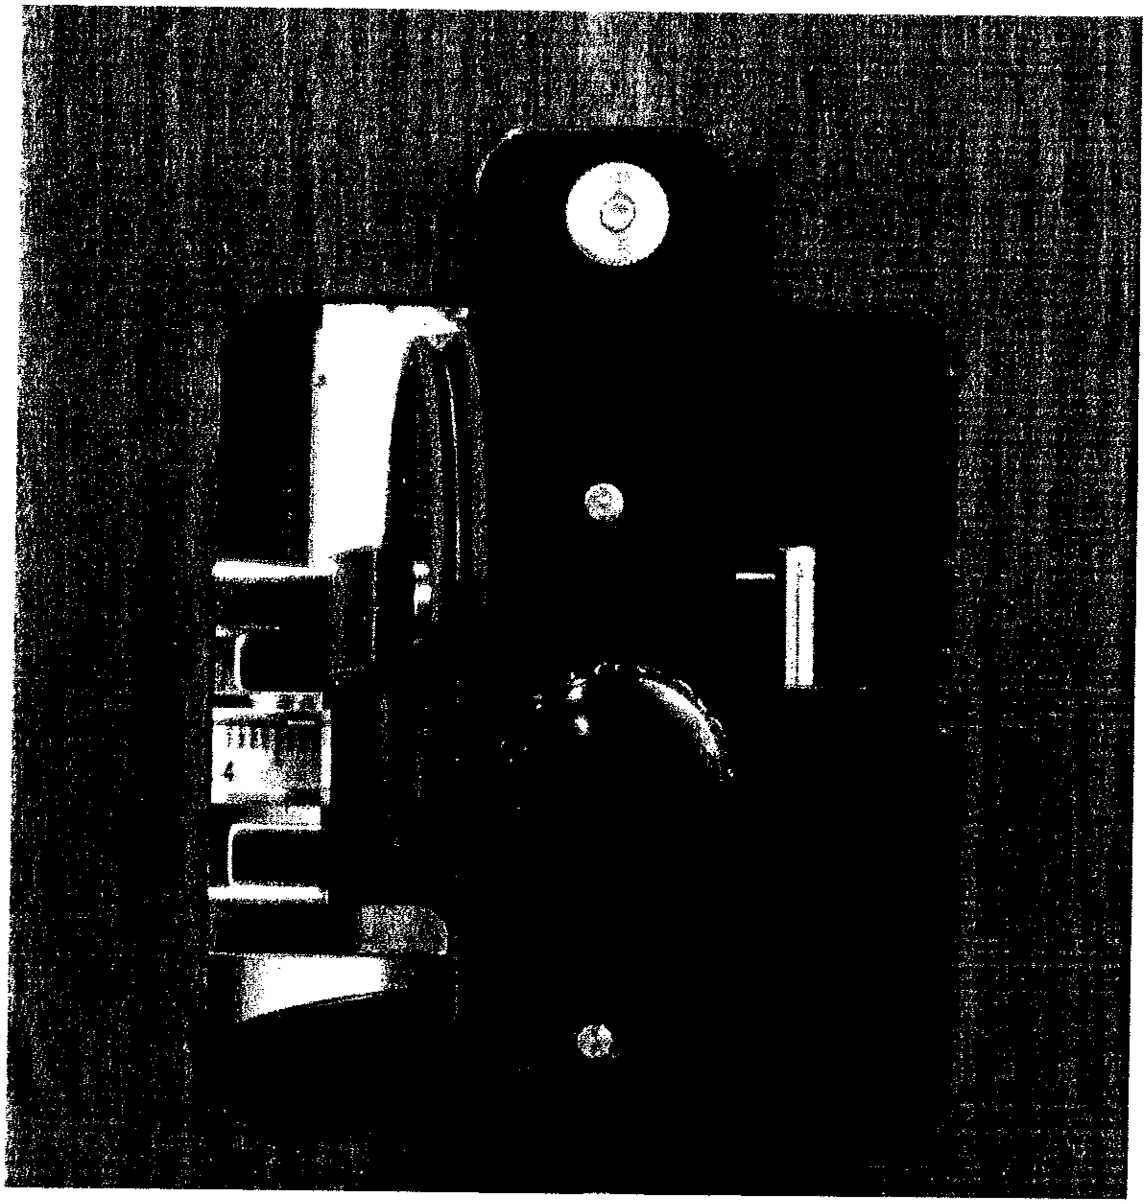

Fig. 1.0a: One of two taut lines exerting a static torque on the rotor

It would also leave an unbalanced sideways force of *mg* acting sideways on the rotor, which would pull the fiber out of a straight line. So for best results, you can arrange for <u>two</u> strings, each with tension, *mg*, each to exert a torque, *Rmg*, for a net <u>torque</u> of *2Rmg*, but (artfully) cancelling the net <u>force</u> on the rotor. See the diagram below for the arrangement that works. In fact, in this arrangement, a longer single piece of line can be used, if it makes a U-turn at the little screw-head in the rotor's hub.

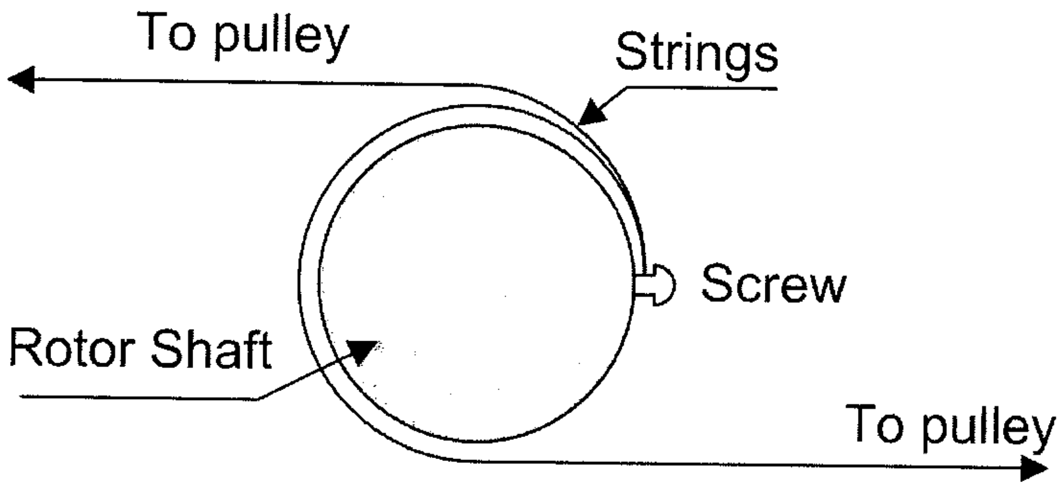

Fig. 1.0b: Using a single line to support both weights

You can arrange for this torque to be positive or negative -- that's the purpose of the other dowel pin on each magnetic-damper structure. In addition to varying the mass m on each side, you can also vary the radius, R, at which the torque is applied, by using either of the black plastic extra hubs that come with your Oscillator. Because they have radial slits, these can be slipped onto the rotor without needing to remove the fiber. Use the two white nylon thumbscrews threading into the rotor shaft to hold the plastic hubs in place.

Note that you should keep the applied torque within bounds, to keep the rotor's angular deflection under 2 radians. **For the thinnest torsion fibers, or the largest hubs, you will *not* want to use the maximal 400 grams on each string.**

1.1 Angular response to static torque: the torsion constant

Section 1.0 will have taught you how to use masses and strings to apply computable torques to the rotor of your Torsional Oscillator. This section shows you how to use the results to find the 'torsion constant' of your torsion fiber. It's the analog of the 'spring constant' of a spring.

You should have some line(s) and masses set up so that the lines are running free, passing smoothly over two pulleys, and pulling purely tangentially on the hub of your rotor. If all is well, your rotor will have turned angularly from its former equilibrium position, and both masses should have descended. If you've released the rotor from rest, and if its motion is undamped, it'll be oscillating away; but it will settle, or you can intervene by hand to help it settle, at a new equilibrium position.

Your independent variable is the torque, $2Rmg$, you've applied, and the assumption is that at equilibrium, the fiber has been twisted enough to be supplying a counter-torque of equal magnitude. Your dependent variable is the angular displacement of the rotor from its original equilibrium position.

For a first investigation of this sort, you'll want to read raw angular positions on the 1.0-5.0 radian scale that's in place on the periphery of your copper rotor disc. (What makes this a radian scale? That is to say, how would you design it so it really gives angles in radians, and not just markings in arbitrary units?) You can use this scale by old-fashioned eyeball methods, sighting through the lines on (both sides of) the sighting plate, and reading the scale that way. The 1-radian divisions are marked with numbers, and 0.10 radian divisions have a major line, while the minor lines mark 0.02-radian divisions. You can even interpolate your readings, perhaps to 1/5 or 1/10 of the smallest divisions, to achieve higher angular resolution.

Such a reading gives you the raw angular position, θ_{raw}. Be sure to make such a reading even before you apply any torques. (That reading might fall near, but not at, the 3.0-mark at the center of the scale.) But after you record θ_{raw} for all your applied torque values, you might want to form the **angular displacement variable,** θ, by subtracting out the θ_{raw}-value that applies for zero external torque.

Now you can graph effect (angular displacement) as a function of cause (static torque applied), and by further deduction you can construct a graph that shows inferred torque exerted by the fiber as a function of the angular displacement of the rotor. The simplest model for the fiber's torque is

$$\tau = -\kappa\,\theta$$

where κ is a constant, of dimensions N·m/rad, which gives the torque per radian of angular displacement. What is your experimentally-measured value of κ?

The torsion constant you've measured can be related to more fundamental parameters, by noting that it arises from the twist of two sections of torsion fiber (one above, one below, the rotor, and each of length 254 mm), and that it arises from twisting a fiber of radius r. If you make this calibration for more than one fiber, you can check a prediction of elasticity theory that κ is proportional to r^4.

1.2 The angular-position transducer

Section 0.2 will have introduced you to the operation of the angular-position transducer on your Torsional Oscillator. It maps the instantaneous angular position of the rotor to a voltage output available on the front panel of the instrument. Its operation depends on the electrical capacitance in the 3-decker 'sandwich' at the bottom of the rotor shaft. This makes it a real-time sensor which works with response time about 10 ms, and without any mechanical contact or friction.

If you're set up to measure the torsion constant of your Oscillator according to section 1.1, you'll have just what it takes to bring the rotor to a series of angular positions whose angular coordinates you know in radians. These are also a set of positions at which you can read the output voltage of the angular-position transducer, so you can calibrate it against a scale in actual radians.

It may be helpful to apply some magnetic damping to your system, to help it settle to a displaced equilibrium position faster. You may use a digital multi-meter, or some more sophisticated tool, to measure the voltage emerging. Plot the effect (transducer output voltage) as a function of the cause (raw angular position, θ_{raw}, or displacement, θ, from equilibrium). Your sensor has been designed to give a very nearly linear variation, with some caveats--

- there's a zero-offset adjust, which is a convenience that you may use -- you can adjust its ten-turn dial at the start to deliver an output of 0.000 DC Volts at the equilibrium position of the rotor.
- there's an issue of the stability of the zero setting -- if you breathe some humid air into the capacitor sandwich, you'll see that this (and other environmental changes) can make a difference.
- as with any transducer, there's the issue of 'noise' -- the output voltage shows fluctuations about its nominal value. If you see fluctuations of a few mV with about 1-second periodicity, these are probably due to real, but very small, actual rotational oscillations of the rotor.
- as with any transducer, there are limits to the range over which its response is linear. For your capacitive transducer, you can expect linear response over a range of ± 1.45 radians, but marked departures from linearity beyond this.
- From a suitable graph, you can write a linear algebraic model that maps angular displacement θ to angular-position output $V_{pos}(\theta)$. That model can be inverted, so that you can go backward from electronically-read V_{pos}s in Volts, to inferred angular position θ in radians.

The considerable sensitivity of your transducer can be used to discover some new physics. The transducer's output gives calibration constant, $\Delta V_{pos} / \Delta\theta$, of about 2 V/rad, and since the stability of its output is better than 2 mV, a displacement of under 1 mrad (1 milli-radian, about 1/20 of a degree) can be detected. Now let the (torque-free) rotor settle to equilibrium, record a V_{pos} reading, displace (by hand) the rotor by +1 rad, let it go back to equilibrium, take another V_{pos} reading, displace the rotor (by hand) by -1 rad, and finally let it go back to equilibrium again, taking a final

V_{pos} reading. You may find the torsion fiber displays a memory of where it's been before -- this is called *hysteresis*, and it's a real effect. You might see if there's a small-displacement regime in which it's absent (or undetectably small), or you might look to see if there's a threshold size of angular displacement that is required to show the onset of this sort of behavior.

1.3 Periods of Oscillation: Modeling rotational inertia

This section will show you what can be deduced from the simplest *dynamic* measurements on your Torsional Oscillator. The measurements make use of the oscillations about equilibrium that you've seen in the motion of your oscillator. You'll need to have worked through the operation of the Oscillator in section 0.2.

You may have had occasion to use the magnetic dampers in your Oscillator, but now you may go to the other limiting case of *minimally* damped motion. You'll have found good ways to hand-excite the torsional oscillations of your system, using the hand-pumping that's instinctive to apply to the extra clamp mounted low on the torsion fiber.

The motion you see, once you're done exciting the system resonantly, is periodic motion, and it's conventional to define T as the period, the duration of one full cycle of the motion. It's surprisingly easy to measure this period quite accurately, even with a stopwatch, provided you realize that timing the duration of (say) 10 full cycles is a better approach than timing one full cycle 10 times. (Why?) Of course you also have an electronic output of the actual position waveform, and you can get more precise results from acquiring and fitting this waveform.

The reason the period, T, is worth measuring is that it is (to a very good approximation) independent of the amplitude of the oscillation, and thus it is characteristic of the oscillating system, and not the conditions of excitation. If you know the prediction for the period for the mass-on-a-spring system,

$$T = 2\pi\sqrt{(m/k)}$$

then you can perhaps believe the analogous result for this kind of rotational motion:

$$T = 2\pi\sqrt{(I/\kappa)}$$

where κ is the torsion constant (previously measured in 1.1), and where I is the rotational inertia (also known as the moment of inertia) of the rotating system.

You can test this prediction using your (statically) measured value of κ, and using an estimated value of I. Use section 6.1 to get the data you'll need to make an estimate for I, assuming I is dominated by the contribution of the copper disc. (Why is this a good assumption?) You can also use the same data table to estimate additions to this I-value for other parts of the rotor structure. With an experimental value of κ, and a good estimate for I, go ahead and predict the period T, and compare with your prediction.

But there's much more you can extract from the result for the period, T. Rather than complete the truly tedious detailed modeling of the rotational inertia of the rotor with its complicated shapes, there's provision in your Oscillator to add very precisely modeled contributions to the inertia. These are provided by the brass quadrants that are stored on the lower shelf of your box. They mount via their dowel pins to the circle of holes machined into the top of your copper disc. Now letting I_0 represent the rotational inertia of the rotor as-is, and letting ΔI represent the added rotational inertia contributed by one brass quadrant, you can write

$$I = I_0 + n\,\Delta I$$

where n is the number of brass quadrants added. Now transform the prediction for T above, to show that $(T/2\pi)^2$ ought to be a linear function of n, and use this result to understand what the slope, and the intercept, of this linear dependence ought to be.

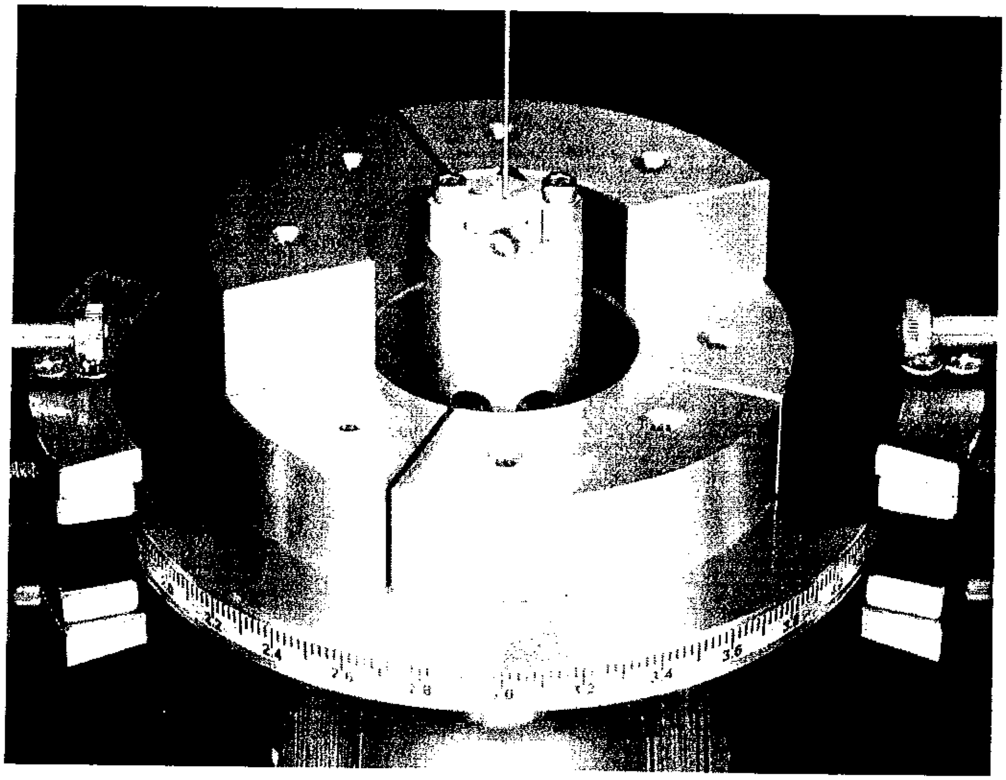

Fig. 1.3: Brass quadrants mounted on the rotor

With that as motivation, go ahead and take data, by your preferred method, for the period of the oscillator as a function of n. You might want to add two brass quadrants at a time (to keep the rotor balanced). The masses stack best if you install them staggered in bricklayer's fashion. Plot your data in the form suggested by the theory, and extract the coefficients of the best linear fit.

To make use of the coefficients of this fit, you'll need to compute ΔI for a brass quadrant, but this is rather easy, since its inner and outer surfaces have been machined so as to be *circular* about the axis of the rotor. With a value for ΔI, you can use the *slope* of your fit to extract a new, and now *dynamically*-measured, value of the torsion constant κ of your torsion fiber. How well does it agree with the value measured by (wholly independent) static methods?

And with that value of κ, you should be able to use the *intercept* of your fit to deduce a value for I_0, the rotational inertia of the unadorned rotor. How well does that value compare with your previous estimates of I_0?

1.4 The angular-velocity transducer

In section 1.2 you've seen that the Torsional Oscillator comes equipped with a real-time analog electronic angular-position sensor, and you've calibrated that sensor. In this section, you'll learn about a completely separate transducer that gives you a real-time indication of angular *velocity*, and you'll learn one way to calibrate it.

Have a look at the rotor shaft of your Oscillator, and notice that in the center of its length, lying at the center of a set of coils, there's a permanent-magnet structure. The stack of four NdFeB magnets visibly moves relative to the coils as the rotor turns. The motion of the magnets with respect to the coils induces an emf in the coils, given by Faraday's Law of Induction. You can use the front-panel toggle switch to send this coil emf to the two output connectors on the front panel. For starters, you'll want to use the right-hand one of those two outputs, as it is better filtered against electronic noise.

The first thing you can check is that you get no output from this connector when the Oscillator is at rest. The next complication comes from the fact that you can't do a static calibration of an output that claims to be a velocity signal! In fact, the calibration constant you want is the number of Volts you'd get out per unit angular velocity of the rotor. The problem is that you can't maintain an angular velocity of (say) 1 radian/second for very long. So instead of trying to calibrate using (a series of) constant angular velocities, it's easier to calibrate using a motion you've already studied -- simple harmonic motion.

Perhaps the best method is to use a 2-channel oscilloscope (or equivalent) for acquiring the angular-position, and the claimed angular-velocity, signals simultaneously. (Set the 'scope inputs to DC coupling.) If you hand-pump the oscillator up to a motion of (say) ±0.1-radian excursion, you'll get an angular position signal, $V_{pos}(t)$, and using your previous calibration of this sensor, you can deduce the angular-position function, $\theta(t)$. Rather than try to differentiate that function numerically, you might instead extract the amplitude, and the frequency, of the signal (by fitting, or some other method), so you can write a mathematical model of the position function. Now you can perform analytic differentiation of the angular-position function to get a prediction for the angular-velocity function.

Compare such a prediction for $d\theta/dt$ to the voltage signal $V_{vel}(t)$ you acquired from the velocity transducer, and see if the two signals have the same shape -- ideally, they should differ only by a scale factor. Find the value of the scale factor that makes the signals agree best, and that's the sensitivity s of your sensor -- that is to say, you've found the value of s such that

$$V_{vel}(t) = s\, d\theta/dt$$

The units of s are Volts/(rad/s) or V·s/rad. Of course, you can invert this relationship, so that you can hereafter convert from a measured $V_{vel}(t)$ to an inferred angular-velocity function, $d\theta/dt$.

1.5 Oscillations, viewed in the 'phase plane'

You now have completed the calibrations of those transducers which will give you instantaneous (voltage) values for the angular-position, and angular-velocity, of your Torsional Oscillator. So far you've viewed these two signals as separate functions of time. Now it's time to get introduced to a valuable presentation, of one of these signals against the other, in a new 'space' called the *phase plane*.

The phase-plane view of the state of your system is best seen in real time on a 2-channel oscilloscope that has XY-capability. You'll want your Oscillator to be nearly undamped, and you'll want to hand-pump it up to oscillations of moderate amplitude. Connect the $V_{pos}(t)$ signal to channel 1 of your 'scope, and get a view of the angular-position signal. Next, use another cable to bring the $V_{vel}(t)$ signal to the channel-2 input of your 'scope, and get of view of it, simultaneous with your view of channel 1. Again, use DC coupling at both inputs, and of course you can adjust the 'scope's sensitivities and zero-offsets to get a nice over-and-under view of the two waveforms.

You should, in this sort of view, be able to see that the position and velocity signals are *not* in phase, nor 180° out of phase, but in a different phase relationship. It's worth learning how to trigger the 'scope on upward-going zero-crossings of the ch. 1 angular-position signal -- this is equivalent to picking an origin of time, or a t=0 point, such that the position waveform can be described by a pure *sine* function. Relative to this choice of origin, what sort of function describes the angular velocity?

From a mathematical model of angular position as a function of time (as a sine function) you should be able to show that you expect the angular-velocity waveform to be a *cosine* function. (If polarities come out that way, it might be a negative-cosine function that you see -- no worries.) So the position and velocity waveforms are both sinusoids, but 90° out of phase.

Now you're finally ready to view these two signals, not as functions of time, but one vs. the other, simply by asking your 'scope for an XY-display of the signals you're already getting. Now the 'scope is in the Etch-a-Sketch™ mode, displaying a point with coordinates (V_{pos}, V_{vel}) which dances around the screen because both coordinates are changing with time. You might want to adjust the scale factors on your 'scope, and you might also want to still your Oscillator, and put the resulting (0,0) point at the center of your screen.

View the trajectory of your moving point, perhaps using the memory or persistence functions on your 'scope to see the path traced out by the point. It's tracing out a 'locus in the phase plane', a fine view of the mathematical trajectory of position and velocity. Now have some fun: still the rotor, and hand-turn it to some non-zero position, but hold it at zero velocity. What do you see in the phase plane? Next, release the system to start its time evolution -- what trajectory do you see? What does the mathematical description of position and velocity lead you to expect for a trajectory? How does the trajectory change if there's some non-zero damping applied to the oscillator?

One of the values of a starting point (V_{pos}, V_{vel}) in the phase plane is that (via some calibration constants) it stands for a particular choice of values (θ, $d\theta/dt$) right at the instant you let go of the system. That is to say, at any instant, the location of the point in the phase plane stands for just the sort of information that a second-order differential equation needs as 'initial conditions'. Since specifying the initial conditions fully *determines* the further time evolution in such a differential equation, it follows that only one trajectory can pass through a given point in the phase plane -- certainly the trajectory can't cross itself at any point.

You can get a look at this kind of 'deterministic time evolution' by testing for repeatability. Set a moderate level of damping, and hand-hold the rotor at a given position, and zero velocity, for a first hold-and-release trajectory, which you can capture on your 'scope view. Now intervene by hand to bring the oscillator's state right back to your original starting conditions -- that is to say, use your real-time 'scope view to come back to the same starting position and velocity you used before. Do a fresh release from this position, and see if it's true that that the system, given the same initial conditions, will exhibit the same post-release solution by retracing the former trajectory with a fresh one lying right atop it.

1.6 From position and velocity to energy

This section assumes that you've completed sections 1.2 and 1.4, so that you have calibrated angular-position and angular-velocity transducers operating in your apparatus. It also assumes that you've completed sections 1.1 and 1.3, so that you have numerical values for the torsion constant, κ, and the rotational inertia, I, of your system. It further requires that you have a data-acquisition system that can collect time records of the two waveforms, $V_{pos}(t)$ and $V_{vel}(t)$, that emerge from them. The payoff is a chance to investigate the detailed behavior in time of the kinetic energy, and the elastic potential energy, of the oscillating system.

Once again, you'll want to excite a moderate-amplitude oscillation of your system, and then let it continue undamped in time. You'll want to acquire data that span at least a few full cycles of the oscillation. What you get is a series of points, probably equally spaced in time, for $V_{pos}(t)$ and $V_{vel}(t)$. Now if you have the two sensors calibrated, you can transform such voltages into the functions that gave rise to them, respectively the angular position, $\theta(t)$, and the angular velocity, $[d\theta/dt](t)$. And those functions are of interest because the energy of the system comes in two forms:

elastic potential energy, $U = (1/2)\,\kappa\,\theta^2$

and

kinetic energy, $K = (1/2)\,I\,[d\theta/dt]^2$

So you can form inferred values of these two quantities for each of the time values for which you have acquired information. You should pause to check that you have your units right, so that each of your energy values comes out in actual Joule units.

Now you can plot the time dependence of U(t) and K(t), and you should see that each of these functions has not one, but two, maxima per cycle of oscillation. In addition, you should be able to say why, and when, these maxima occur.

Finally, you can form, and plot as a function of time, the additional quantity

$K(t) + U(t)$

which is called the mechanical energy of the system. It should be a *constant* for an undamped oscillation, even though both of its pieces undergo full-scale oscillation. What sets the value of that constant? Can you get a data set displaying a *different* value for that constant?

If you succeed in this data-acquisition and data-transformation process, you might want to see what happens in the case of a *damped* oscillator, in which mechanical energy is *not* expected to be conserved.

NATIONAL INSTRUMENTS™

Certificate of Calibration

Board Information		Certificate Information	
Part Number:	187570C-01	Certificate Number:	26601
Description:	PCI-6023E	Date Printed:	10-AUG-2002
Serial Number:	101817E	Certificate Part Number:	184632A-01
Calibration Date:	10-AUG-2002	Ambient Temperature:	22 °C
Calibration Interval:	12 months	Relative Humidity:	52 %
Calibration Due:	10-AUG-2003		

National Instruments certifies that at the time of manufacture, the above product was calibrated in accordance with the applicable National Instruments procedures. These procedures are in compliance with relevant clauses of ISO 9002 and are designed to assure that the product listed above meets or exceeds National Instruments specifications.

National Instruments further certifies that the measurements standards

1.7 Projects

Here are some projects that make use of the skills and calibrations you've acquired so far.

1.7.1 Dependence of the torsion constant κ on the *diameter* of the torsion fiber

There's an initially surprising prediction of elasticity theory that the torsion constant you'll get from a fiber of circular cross section depends on the *fourth* power of its diameter (other things being equal). You can confirm this variation using the Torsional Oscillator, either by having four units set up with distinct fibers, or by changing fibers in a single unit.

Section 0.4 of this manual takes you through the detailed procedure for removing and replacing a fiber. It's not a very fast or simple procedure, and it's not the only way to test this fourth-power dependence -- you can perhaps devise other experiments that will make the same test.

If you want more fibers to test, try your local hobby shop as a possible source for music wire, also known as 'piano wire'. Of course you can also try other materials: you may have copper wire on hand for electrical purposes (often called 'magnet wire'), and you can try stainless-steel wire too.

1.7.2 Dependence of the torsion constant κ on the *tension* in the torsion fiber

Since the Torsional Oscillator has in place a way to vary the tension in the fiber, and since the 'plucking test' reveals that it can in fact be changed over some range, you might wonder what effect this has on the torsion constant of the tautened fiber. The results are surprising enough, and easy enough to obtain, that you might want to test for this dependence.

You can choose your favorite method for measuring either absolute, or relative, values of the torsion constants (see sections 1.1 and 1.3) That's the dependent variable in this investigation, but the independent variable is the level of tension in the fiber. It's easy to vary that with the tensioning bar at the top of the apparatus, but it's a bit harder to measure it directly. So we've built into the Torsional Oscillator a method for *indirectly* determining the tension in a fiber.

That method depends on the frequency of oscillation (psychologically, related to the pitch of the audible sound) of the guitar-string modes of the fiber's upper or lower segments. You'll need to look up, or derive, the expected relationship between this frequency of transverse vibrations, the length of the segment of wire, its mass per unit length, and the tension in it. The value of this relationship is that the upper and lower segments of torsion fiber in your Oscillator both have a free length very near 254 mm, and the fibers involved also have a well-known mass per unit length (see section 6.1 for details). So if you can determine the frequency of the musical note you can hear by plucking, you can infer the tension to adequate precision.

To determine that frequency, the easiest method is to recruit a helper with the gift of 'perfect pitch'. But lacking such a useful co-worker, you can get the frequency electronically. Built into the electronics board that reads out the angular-position signal is a little microphone/amplifier combination, and its output is brought to a test point actually accessible on the front the printed-circuit board in question.

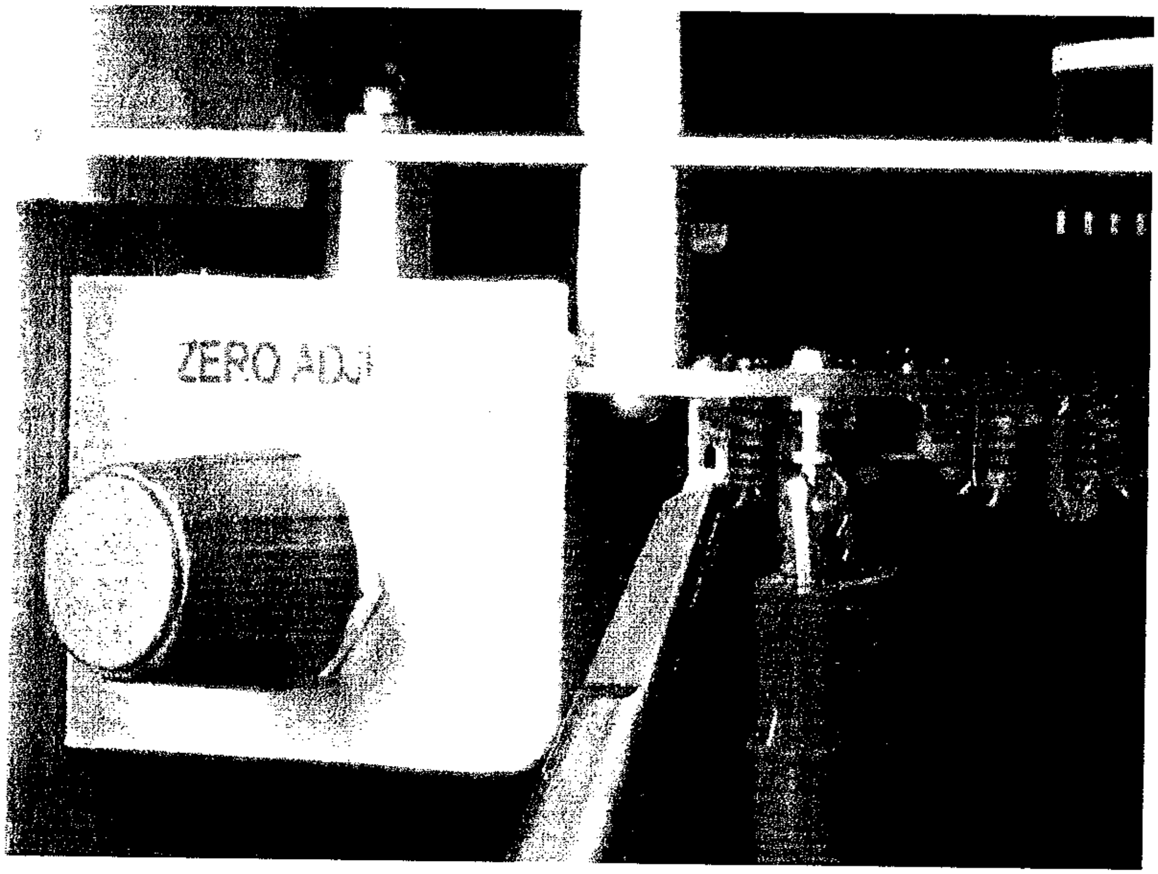

Fig. 1.7.2: The signal and ground points for picking up the audio signal

Remove the wire-clamp structure ordinarily in place (for hand-exciting torsional motion) on the lower segment of the fiber. Test to hear the pure musical note you get when you pluck the center of the fiber's lower section. Now use a suitable tool to get an electronic version of that sound, and use either the waveform in time, or its Fourier transform, to infer the fundamental frequency that's present. (There will be harmonics of this fundamental also present, especially right after plucking the fiber.)

You can safely vary the tension in the fiber over a considerable range, perhaps up to a few hundred Newtons. At some level, the fiber will slip in its clamps, and the plucking test will tell you when that starts to happen. You are scarcely likely to break the fiber by mere pulling -- even the thinnest of the fibers supplied has a nominal breaking strength over 1000 Newtons. It's no accident that the cables of suspension bridges are built using the same kind of steel fibers as the music wire that you're using.

1.7.3 Dependence of the period, T, on the *amplitude* of oscillation

You've found a way to measure the period of torsional oscillations in your apparatus, and you've seen the period, T, varies systematically with the rotational inertia of your rotor structure. But you've implicitly assumed, or perhaps expressly tested, the claim that the period of oscillation is independent of the conditions of excitation, and in particular is independent of the *amplitude*, of the oscillation.

For reference, you should know that the historically important case of the period of a pendulum does *not* have this 'isochronous' property. For an ideal (point-mass) simple pendulum, the usual prediction that the period

$$T = 2\pi\sqrt{(l/g)}$$

is in fact a small-angle approximation, and for non-negligible amplitude of oscillation, θ_m, there are corrections in the form of a power series

$$T(\theta_m) = T(\theta\approx 0)\ [\ 1 + (1/16)\ \theta_m{}^2 + \ldots\ ]$$

These corrections are *not* small, at the level of precision of interest in pendulum clocks. The question is: do similar small-angle approximations exist in your torsional oscillator? Or to put it another way -- if you measure $T(\theta_m)$ empirically, and fit it to an equation of the form suggested by the pendulum result, what value, or what upper bound, do you get for the coefficient of the $\theta_m{}^2$-term? If that coefficient is as large as 1/16, your oscillator is no more isochronous than a simple pendulum. If you can show the coefficient is smaller, then a 'torsional pendulum' clock might be free of one source of time-keeping imprecision that bedevils a pendulum clock.

1.7.4 Alternatives to the brass quadrants -- another inertia calculation

In section 1.3 you had occasion to use the 'brass quadrants' as a way to vary systematically the rotational inertia I of your oscillating system, and you used computed values of the rotational inertia contribution, ΔI, that each quadrant provides to deduce other parameters of your system. Now here's another and alternative way to add inertia to your rotor, and to test predicted values of rotational inertia contributions.

Your Torsional Oscillator comes with eight precisely-crafted steel spheres, better known as ball bearings, which can be mounted on the same circle of holes in your copper rotor disc as you've used to locate the dowel pins on the bottoms of the brass quadrants. The little conical depressions at the tops of these holes provide a simple way to locate precisely the centers of those spheres on a circle of well-known radius -- Section 6.1 for dimensions. (Note that you can put all eight spheres onto to copper disc at once, but that you can't use arbitrarily large amplitudes of rotational motion and still expect the spheres to stay in place.) So just as the brass quadrants each provide a ΔI_{brass} that changes the period of your oscillator, so each steel sphere provide a ΔI_{steel}, and this ΔI will also show up in affecting the period of torsional oscillations of your system.

A procedure similar to that of section 1.3 will give you period data that can be used to check two results: can you get an I_0 value using steel spheres that matches the result you get for brass quadrants? And can you get, from a suitable graph of your data, a good empirical value for the ratio, $\Delta I_{brass}/\Delta I_{steel}$?

Of course you can measure the masses, and the relevant dimensions, of the brass and steel objects, and now the goal is to compare the empirically measured value of the ratio, $\Delta I_{brass}/\Delta I_{steel}$, with the value computed theoretically from the masses and dimensions. In fact, the real motivation is to see if you can make the empirical measurement accurate enough to resolve the difference you get in that predicted ratio, depending on whether you treat the steel spheres as point masses, or more correctly as extended objects.

1.7.5 Larger-angle behavior of the velocity transducer

If you've viewed the angular-position transducer output for oscillations at various amplitudes, you've seen sinusoids for all amplitudes (up to the linearity limits of that transducer). But if you look at the output of the angular-*velocity* transducer, for the very same oscillations, you'll see the shape of the waveform systematically changing -- from sinusoids for low amplitudes, to something that looks more like a triangle wave at intermediate amplitudes, and then gets stranger still at large amplitudes. What's going on here?

The first thing to understand is that for any amplitude, the velocity-output is still periodic, and of the same period as the position-output. So what you're seeing is a periodic waveform, but one that's *not* a simple, single, sinusoid. The velocity-transducer output is showing what in the audio world would be called 'harmonic distortion'.

But what's the physical explanation of this? It's not 'distortion' in some electronic processing chain, as it would be in the hi-fi case. In fact, there are no active devices at all in the velocity-transducer -- you can check that the output looks identical whether line power is supplied to the unit or not! The explanation has to be sought in the physical mechanism by which the velocity signal is formed to begin with: that's Faraday's Law of Induction.

Consider the Helmholtz coil as a pick-up coil, and the permanent magnets on the rotor as a source of magnetic flux, Φ, that passes through the pickup coil. The emf that provides the velocity-transducer signal is related to $d\Phi/dt$ by Faraday's Law. But Φ is changing with time for a double reason: Φ depends on the angular position θ of the magnets relative to the coil, and θ itself is changing with time. So by the chain rule we can write

$$d\Phi/dt = (d\Phi/d\theta)\,(d\theta/dt)$$

and the last factor is indeed the angular-velocity terms we want. What is the first factor?

We model the permanent magnets as the source of a dipole field, and we assume that the $\theta=0$ position of the rotor has the magnets' magnetic moment perpendicular to the axis of the Helmholtz coils. Then it is a very good model to assume that the flux behaves as

$$\Phi(\theta) = \Phi_m \sin(\theta)$$

where Φ_m is the maximal possible flux, and where that maximum is only attained when θ reaches 90°, i.e., when the magnets become fully aligned with the coils' axis.

You can test the consequences of this claim quite directly. Suppose the actual mechanical oscillation is a pure and undistorted simple sinusoid, of amplitude, A, and period, T, and suppose we pick t=0 to coincide with a zero-crossing of the angular signal. Then the angular position is given by

$$\theta(t) = A \sin(2\pi\, t/T)$$

and the flux model above predicts that the induced emf, and hence the velocity-transducer signal, will be proportional to

$$d\Phi/dt = (d\Phi/d\theta)\,(d\theta/dt) = \Phi_m \cos(\theta)\,(d\theta/dt)$$

$$= \Phi_m \cos[A \sin(2\pi\, t/T)]\, A \cos[2\pi\, t/T]\,(2\pi/T)$$

This is a complicated function! The second cosine factor has the shape of the expected velocity signal, but it's the first cosine factor that accounts for the complicated waveform that you've seen. For amplitudes, A, that are small enough (A << 1 radian)), this factor can be approximated by 1 for any and all t-values, and the emf is predicted to have the shape of the actual velocity waveform. But for A of size 1 radian, this first cosine factor varies away from a value of 1, to a value smaller than 1, in fact as small as $\cos(\pm 1) \approx 0.5$.

Graph the predicted emf function for several values of amplitude A, and observe the predicted velocity-transducer waveform display some of the same shapes that you have observed experimentally. If you understand this 'forward process', going from assumed amplitude to predicted waveform, you can try the harder 'inverse problem': from a velocity-transducer output recorded by your partner, and acquired with an oscillation of amplitude unknown to you, see if you can find a procedure to deduce, from velocity-transducer output waveform alone, what amplitude of oscillation must have been used to produce the data you're analyzing. So the velocity-transducer waveform contains, in this 'distortion', a key to the amplitude of the motion, which represents amplitude information independent of the use of the angular-position transducer.

Please note also that this physical mechanism that 'distorts' the velocity transducer's output does not distort everything. Show from the mathematics above that it's predicted *not* to distort the implied period of the motion, nor to distort the instantaneous velocity signals you get at those maximal-velocity instants when the angular position is passing through zero.

1.7.6 Filtering, and the angular-velocity transducer

You've seen the angular-velocity transducer in action in section 1.4, and this section examines the signals visible at two distinct outputs on your Oscillator's front panel. This exercise will show you something about *filtering* in a real-life application.

The Oscillator's front panel shows how the coils are connected for velocity-transducer action. *Any* change in magnetic flux through the coils induces an emf in them. This includes flux changes due to the rotor's motion, flux changes due to stray line-frequency fields, and also changes due to all broadcast radio-frequency fields.

High frequencies are partially filtered out by the combination of a 330 Ω resistor and a 10 nF capacitor. This combination defines a low-pass filter of 'time constant',

$$\tau_1 = (330\ \Omega)(10\ \text{nF}) = 3.3\ \mu\text{s}$$

and a 'corner frequency',

$$f_1 = (2\pi\tau_1)^{-1} = 50\ \text{kHz}$$

So frequencies above 50 kHz begin to be suppressed by this filtering action, and this filtered version of the coil emf is always available at the left of the two angular-velocity outputs.

With the rotor at rest, connect a 'scope channel to this left-hand BNC output, you'll see how much 'noise', i.e., unwanted signal not related to the rotor's motion, still appears here.

To reject some of that noise, you can use the right-hand BNC output, which is further filtered by another low-pass filter built of a 10 kΩ resistor and a 1.0 μF capacitor. This filter defines a time constant,

$$\tau_2 = (10\ \text{k}\Omega)(1.0\ \mu\text{F}) = 10\ \text{ms}$$

and has a corner frequency,

$$f_2 = (2\pi\tau_2)^{-1} = 16\ \text{Hz}$$

So the emfs due to the rotor's angular velocity, occurring as they do at 1 to 2 Hz, are fully 'passed' by this low-pass filter, but line frequencies of 50-60 Hz are partially, and higher frequencies are more fully, suppressed.

To see the effect, send both the velocity outputs to the two channels of a dual-trace oscilloscope simultaneously. Now hand-excite the oscillator, and view the two traces to see the velocity signal appear in both channels. View the degree to which the more-filtered signal is less noisy than the other, rawer, signal.

This filtering comes at a cost. The second-stage RC filter not only suppresses frequency content above about 16 Hz, is also time-delays, by about $\tau_2 = 10$ ms, all the low-frequency signals that it does pass. See if you can establish that the existence of that time delay, by comparing your less-filtered and more-filtered versions of the velocity signal.

1.7.7 Time delays, phase shifts, and the phase plane

You've seen in section 1.5 the 'phase plane' depiction of the instantaneous state of the Oscillator, and you've seen the curves in the phase place traced out in time by the 'system point', (θ, $d\theta/dt$). This section takes into account the existence of time delays, or phase shifts, in the two signal channels that you've been depicting in that phase plane.

Let the (undamped) oscillations of the system be modeled as producing a position-transducer waveform appropriate to a position coordinate

$$\theta(t) = A \cos(\omega t + \phi)$$

so that the actual angular velocity is given by

$$d\theta/dt = - A \omega \sin(\omega t + \phi)$$

What you view are two voltages, $V_{pos}(t)$ and $V_{vel}(t)$, which are respectively proportional to these two coordinates' values. Suitably scaled versions of $V_{pos}(t)$ and $V_{vel}(t)$ will fall on the locus of an ellipse in the phase plane, because there are constants, *a* and *b*, such that

$$[V_{pos}(t)/a]^2 + [V_{vel}(t)/b]^2 = 1$$

(that works out because $\cos^2\theta + \sin^2\theta = 1$). This works perfectly so long as the actual $V_{pos}(t)$ and $V_{vel}(t)$ signals are in fact of cosine and sine character, i.e., 90° out of phase.

But in practice, it's easy for these signals to be phase-shifted a bit *away* from this 90° phase condition. For example, your actual $V_{pos}(t)$ is itself a filtered version of the raw angular-position signal, which imposes a time delay of order 10 ms on the raw signal. Similarly, the right-hand $V_{vel}(t)$ output is also phase-shifted by filtering, and thus also time-delayed by about 10 ms.

If these time delays were to match perfectly, you could still hope for the ellipse in the phase-plane locus to be perfect. But what if the phase shifts don't quite match? You might try a phase-plane plot of $V_{pos}(t)$, and the *less*-filtered version of $V_{vel}(t)$, to see if you can spot any such effects. They'll be most prominent in the case of higher-frequency rotor motion, i.e., with the use of the thickest torsion fibers, and the rotor as light as possible. If you can see a 'tilted ellipse' in the phase plane, you might work analytically on the case of two sinusoids, of equal frequency, that are not exactly 90° out of phase. See if you can show that they still create an elliptical locus, but one who's major and minor axes are not parallel to the coordinate axes in the phase plane. You want to find constants A, B, and C, such that

$$A[V_{pos}(t)]^2 + B\,[V_{pos}(t)][V_{vel}(t)] + C[V_{vel}(t)]^2 = 1$$

and then realize that the constancy of this 'quadratic form' in your two variables still defines an elliptical locus, but one that's been rotated relative to the coordinate axes.

2 Magnetic Torque

You've had occasion to apply torque to the rotor of your Torsional Oscillator using taut strings and gravitational forces, but this section will introduce you to another, non-contact, and electrical way to apply torques to the rotor. The motivations range from some fundamental studies of magnetic interactions, to a very practical way to apply torques that have some chosen, and externally-variable, time dependence.

2.0 Applying magnetic torque

The physical mechanism that causes these new torques is the interaction between the magnetic fields created by the Helmholtz coils with the magnetic moment of the permanent magnets on the rotor shaft. So in these investigations, those Helmholtz coils are no longer serving passively as before, as a pick-up for emfs generated by Faraday's Law. Instead, you'll use the Helmholtz coils *actively*, sending a steady (or time-varying) current through them from some external power source.

To do this, you can flip the toggle switch on the instrument's front panel, to bring the coil connections out to the two grey banana-plug terminals indicated by the schematic drawing. Now you can hook the power supply of your choice to these terminals, and the switch will conduct the externally-generated current through (both of) the Helmholtz coils (the two coils are wired in series). Each coil separately generates a magnetic field, and the two coils together generate a magnetic field which is crafted to be spatially very uniform in the vicinity of the permanent magnets on the rotor.

The series-connected coils have a resistance of 7-8 Ohms, and they can be used with steady currents up to 2 Amperes, (A), or more briefly with currents up to 3 A. The limits are due to thermal dissipation -- at 2 A, the coils will warm up rather slowly, and 3 A they will heat up rather quickly. There's a self-resetting 'fuse' in series with the coils to protect them against overheating -- if you spend enough time running at more than 2 A, the fuse will 'trip', disconnecting the coils. Turning off the applied voltage and waiting a while will allow the 'fuse' to automatically reset, thus allowing continued operation.

There may be occasions on which you want to measure the coil current by other means than an ammeter. In these cases, you can route the coil current, using a jumper wire, through the internal 1.0-Ohm, 1%, 10-Watt resistor inside the instrument, and then monitor the potential difference or 'voltage drop' across this resistor. Now you have available a voltage *surrogate* for the current, with a scale factor of 1 Volt out per 1 Amp of coil current. (This method is better than monitoring the potential difference across the coils themselves -- the coils' resistance will vary with their temperature, which changes during operation.)

2.1 The 'torque balance'

Here's a way to observe magnetic torques quantitatively, by using them to 'balance out' gravitational torques of the sort you used in section 1.1. The goal is to get a numerical measure of magnetic torques, by comparing them against gravitational torques you know how to compute. The procedure will require a variable DC power supply capable of 2 or 3 A output, and it'll need to have a potential difference up to 15 to 25 Volts DC maximum. You'll also want an external ammeter to measure the size of the DC current you're supplying. Finally, you may want to *avoid* the use of the thinnest of the torsion fibers in doing this experiment -- or if you do have the thinnest fiber in place, you'll want to stay with rather modest sizes of gravitational and magnetic torques.

So, set up your oscillator free of any torques, but perhaps with a modest level of magnetic damping to let it stabilize more quickly. Note the equilibrium position, both on the angular scale and the position-transducer output. Now arrange, with the strings, pulleys, hangers, and masses as formerly used, to apply some gravitational torque to the rotor. It will of course move, and eventually settle into some new, displaced, equilibrium position. What's new is the chance to apply a *counter*-torque, magnetically, chosen to bring the rotor *back* to its original equilibrium position. So connect your external power supply to the coils, and find the sign, and the magnitude, of the current that is required to bring the rotor back to its original equilibrium position.

Repeat this 'torque balancing' for a variety of sizes, and both signs, of gravitational torque. As you reverse the signs of the gravitational torque (say by using the strings to pull clockwise rather than counter-clockwise), you'll have to reverse the sign of the current too (say by reversing the connections to the power supply). Form a table, and make a graph, that gives electric current required to balance out gravitational torques of known sizes, and draw a conclusion from this data. From your deduction and your data, make another graph that shows how much torque you're generating magnetically (in usual torque units, of N·m) as a function of how much current you're sending into the coils (in usual current units, of Amperes). See if you can explain why the graph has the shape that you see, and deduce the value of a useful constant, the 'torque per unit current' that your system can generate.

You may have noticed that you haven't needed, in this investigation, to use the numerical value of the fiber's torsion constant. So what difference does the torsion constant make? That is to say, what would guide you in choosing a fiber to install in this apparatus, if you were starting from scratch?

2.2 Angular response to magnetic torque

In section 2.0 you learned how to apply magnetic torque to your Oscillator's rotor, and in section 2.1 you learned to quantify this torque by balancing it against gravitationally-generated torques. In this section, you'll omit the gravitational torques, and just apply magnetic torques to the rotor as an independent variable. The angular-position response of the rotor will be your dependent variable. In the process, you'll learn more about the mechanism by which the Helmholtz coils apply a torque on the rotor's permanent magnets.

Again, start with a rotor free of external torques, and this time, it's 'no strings attached'. Again, you may want to use some damping, so that the rotor will settle at new equilibrium positions smoothly and rapidly. Now your independent variable is the size of the electric current you're sending through the Helmholtz coils, and the dependent variable is the equilibrium angular displacement of the rotor. You may measure the raw angular position of the rotor quite directly using the radian angular scale, or indirectly using the angular-position transducer.

Start with small values of the current, and remember to take data using both signs of current. You might want to plot the data as you take it, since you're going to see initially straightforward data take on a surprising form as you proceed to currents of larger size. Note that you can take data at your leisure up to currents of 2 A or so, but that you'll need to take data rather more quickly for larger currents -- else the self-resetting fuse will shut down the current to protect the Helmholtz coils against overheating.

Now what you see directly is the angular position of the rotor, and what you can infer from equilibrium is that the rotor twists, in response to magnetic torque, until the torque developed by the fiber due to its twist balances out the magnetic torque. Since you've previously modeled the fiber's torque in terms of a torsion constant, κ, you can re-cast your data in the form of magnetic torque achieved, as a function of external current sent into the coil. Make such a plot, displaying torque as a function of current.

Find the regime in which you can model the system as delivering torque proportional to the current, and find the 'torque per unit current' that characterizes this regime. Compare your result with the value you got from the 'torque balance' method of section 2.1. Go on to think about the *other* regime you've now discovered, in which the torque seems <u>not</u> to be proportional to the current. Why is this so? How can it be so, given that in section 2.1 you showed experimentally that magnetic torque <u>is</u> proportional to the current involved?

2.3 Modeling magnetic torque

If you've done the experiments of section 2.1, you've balanced magnetic torques against gravitational ones, and if you've done the experiments of section 2.2, you've balanced magnetic torques against elastic torques due to your torsion fiber. If you've gotten apparently contradictory results, you're ready to resolve this discrepancy by learning something new about magnetic torques.

Start with your Helmholtz coils, which are the source of a magnetic field. Its magnitude *is* proportional to the current you send through the coils, and its direction is along the axis of symmetry of the set of coils. Now turn your eye to the permanent-magnet stack on the rotor, which sits (by design) right in the sweet spot at the center of the Helmholtz coils, and is thus immersed in the magnetic field they generate. That collection of permanent magnets can be characterized by a vector, μ, called its 'magnetic moment'. The magnitude of μ is a measure of the strength of the permanent magnets, and the direction of μ is along the axis of magnetization of the magnets -- at the rotor's initial equilibrium position, that vector points out toward you. Notice that the direction of the magnetic field, **B**, and the direction of the magnetic moment, μ, start out (nearly) perpendicular -- but that they *depart* from this perpendicular condition if and when the rotor takes on a different angular position.

If your apparatus is aligned such that μ and **B** start out perpendicular, then the angle between them in general can be written as

$$\text{angle } \varphi \text{ (between } \mu \text{ and } \mathbf{B}) = 90° - \theta$$

where θ is the angular displacement of the rotor, and where we take θ to be positive in the direction that makes μ and **B** more nearly parallel.

Now you're ready to understand the otherwise very formal definition of the magnetic torque, τ, that a magnetic moment, μ, experiences when it's in a magnetic field, **B**. That claim is that

$$\tau = \mu \times \mathbf{B},$$

a vector cross product.

Show that according to this prediction, the direction of the torque vector is always vertical -- what does that do to your rotor? And work out the predicted magnitude of that magnetic torque, and show that it can be written as

$$\tau = \mu\, B \sin \varphi = \mu\, B \sin (90° - \theta) = \mu\, B \cos \theta$$

Now write the equation that describes the balance between the elastic torque of the fiber, and this size of magnetic torque, and see if you can understand that it predicts that not θ, but rather $\theta/\cos \theta$, ought to be linear in the coil current that you used in section 2.2. And now re-plot the data of section 2.2 to display the experimental behavior of $\theta/\cos\theta$ as a function of the coil current, i. Success in your new plot stands for validation of the form of the angular dependence predicted by the vector-product torque model discussed above.

2.4 The coil constant of the Helmholtz coils

This section shows you how to compute the 'coil constant' for the Helmholtz-coil pair, which is the number that gives the magnetic field at its center per unit current through its windings. This number can be deduced from fundamental constants and measured dimensions, and it's valuable for extracting the size of the magnetic moment of the permanent magnets with which it interacts.

First of all, the dimensional parameters. The two coil bobbins forming the pair each carry 201 ± 1 turns of #22 AWG copper wire, and the two 201-turn coils are wired in series (so their fields *reinforce* at the geometrical center of the assembly). As assembled, the coils-in-series have a (d.c.) resistance near 7.6 Ω, and a (low-frequency) inductance of 17 ± 2 mH.

The coils are wound on bobbins of outer diameter, 4.90", with a groove of depth, 0.45", and width, 0.40". (Here 1" means 1 inch, defined to be 0.0254 m exactly.) The windings of wire thus start with inner diameter, 4.00", and end near a nominal outer diameter, 4.80". The two bobbins are mounted on a structure that gives the grooves in the bobbins a center-to-center separation of 2.20". This also matches the radius of the 'typical turn' of wire, which achieves the Helmholtz design optimizing the field uniformity near the center of the assembly.

The simplest model of the coils is to suppose that all 201 turns on each coil somehow have collapsed right at the center of the winding space on the bobbins, and they thus all form circles of radius, a = 2.20", in two planes at locations, $z = \pm$ 1.10", measured along the common axis of the two coils. Now finally the field at the center can be computed from a result itself derived from the Biot-Savart Law, which gives the field on the axis of a single-turn circular coil of radius, a, at a test point lying distance, z, above (or below) the plane of the coil:

$$B(z) = (\mu_0 \, i \, a^2/2)\,(a^2 + z^2)^{-3/2}$$

where $\mu_0 \equiv 4\pi \times 10^{-7}$ T.m/A in the SI system of units. In the case of a Helmholtz coil system with N turns on each of two coils, this gives for the field at the geometrical center the result

$$B(\text{center point}) = (\mu_0 \, N \, i \, a^2)\,(a^2 + (a/2)^2)^{-3/2}$$

$$= (\mu_0 \, N \, i \, /a) \times 8/(5\sqrt{5})$$

and finally the coil constant desired is given by the ratio, B/i.

This result can be evaluated to express B/i in Tesla/Ampere, or more conveniently in mT/A, and it is a fine first model for the coil constant. In actual fact, the N=201 turns are spread out in space, having differing a and z values, so more detailed modeling is possible. The easiest models consider pairs of circular turns, (one on each bobbin) having distinct a_j- and z_j-values, using the Biot-Savart result for a single pair of circular coils, and then summing over j the 201 terms to yield the result for the full coil.

2.5 Deducing a magnetic moment

If you have worked through sections 2.1 – 2.3, you have measured and modeled the torque due to a magnetic moment, μ, immersed in a magnetic field, **B**, and if you have a result from section 2.4, you have a proportionality constant between the current, i, you apply to the coils and the field, B, that they produce. This information can know be combined, to give a numerical result for the magnetic moment, μ.

So to combine results from earlier sections, realize that for the case of a moment perpendicular to a field,

(torque/unit current) = (torque/unit field) x (field/unit current)

and you can now put in the values you've measured to give a numerical result to the (torque/unit field). But for the conditions specified, that quotient is in fact the numerical value of the magnetic moment, μ.

Emerging from this definition are the units μ, which are (N·m/rad)/Tesla, or for short, J/T. You should be able to show that the units of μ are also given by $A \cdot m^2$, a product of current times area.

You have now measured an actual magnetic moment in SI units by a rather fundamental means, sometimes called 'torque magnetometry'. The numerical value of μ you've extracted is by itself perhaps not too interesting, but you can put it into context by performing two calculations.

First, you can ascribe that magnetic moment to the full volume of the four discs of NdFeB that make up the permanent magnet on the rotor. Since section 6.1 gives the dimensions of the magnets, you can compute the volume of space occupied by the magnets, and interpret it as the volume in space through which the magnetic moment is spread. Now you can get an actual numerical value for the magnitude of the otherwise highly abstract vector field, **M**, defined in electromagnetic theory as the volume density of magnetic moment, or magnetic-moment per unit volume. It has units of $(A \cdot m^2)/m^3$, or A/m.

To be even more concrete, you can now compute the numerical value of the magnetic moment of each of the four individual discs of NdFeB making up your permanent-magnet stack. The external magnetic interactions of that disc are just the same as if the disc was made of (say) wood, but it had a one-turn ribbon of current flowing around the periphery of the disc. To find the magnitude of this 'Amperean current', you can see that a one-turn loop of current, i_A, flowing around the rim of a disc of radius, r, would give it a magnetic moment of $i_A \pi r^2$. If you equate this number to the moment you've deduced for one disc, you can see how enormous the effective current, i_A, really is. Considering that it flows (or at least, acts *as if* it flows), with no dissipation at all, you can see in some numerical sense how remarkable a piece of material a modern permanent magnet really is.

2.6 Reciprocity: two views of a magnet-in-coil system

You have now seen the magnets on the rotor shaft interacting with the Helmholtz coils in two quite distinct ways. In section 1.4, you saw the magnets as active, the coils as passive, and the system act like a generator, producing an emf proportional to the angular velocity of the rotor. In section 2.3, you saw the coils as active devices and the magnets as passive, with the system acting a bit like a 'motor', producing a torque proportional to the current you put into the coils. In this section (not required for any future development), you will see that these two views are in fact very closely related indeed, and that the constants you found are not independent.

For simplicity, think of instants of time in which the magnets' axis lie perpendicular to the coils' axis, so that there are no angular complications. Look back to 1.4, and find the numerical value of the angular-velocity calibration constant you deduced there. It will come with units, of the form, Volts/(rad/s), or just V·s/rad. But now think of section 2.3, where you deduced, for the same perpendicular geometry, the torque per unit current in the coils. That number also comes with units, and they are (N·m)/Ampere. Your *first* surprise will come from showing that these two combinations of units, one appropriate to the 'generator' function, and the other to the 'motor' function, of the magnet-in-coil system, are in fact equivalent.

The *second* surprise you should find is that the numerical values of these two constants, found in wholly different ways but applying to the same coil/magnet combination, are nevertheless the <u>same</u> (or at least, equal within uncertainties). This is no accident, and it's called a 'reciprocity principle' -- it turns out to one example of a whole class of relationships that have something of the status of Newton's Third Law. As in that law, the situation has the same two actors, but involves two distinct situations: in one case, object C acts as an active agent on passive victim M, while in the companion case, object M acts as the agent while C is the victim. (Put in 'coil' for C, and 'magnet' for M, and see which view is the 'generator' use, and which is the 'motor' use, of the magnet-in-coil system.)

The proof of this reciprocity principle is quite a bit harder, and in its simplest form is very closely related to the proof of the equality of mutual inductances, M_{12} and $M_{21,}$ for two rigid coil systems labeled #1 and #2. In the case as hand, you can think of the Helmholtz coils as system #1, and the Amperean currents flowing around the periphery of the permanent magnets as system #2, and you can work out quantities like the (flux through #1) per (unit current in #2), and (flux through #2) per (unit current in #1). It's a bit harder to relate these flux quotients to torque and angular displacement, but that will complete the proof.

If you adopt the other viewpoint of accepting a theorist's word of the proof's validity, you have a valuable consistency check on a host of separate calculations in comparing two values as you did above. You also can let the more accurately measured result, perhaps torque per unit current, play a second role, as giving a more accurate value to the less-well-measured emf per unit angular velocity. You even have the advantage of understanding the reciprocity principle that makes vibrating-sample magnetometry possible. In this important instrumental technique, instead of measuring force per unit current, the apparatus measures emf per unit velocity, for a sample acting as a current loop, immersed in a tailored magnetic field.

2.7 Projects

There is a rich collection of projects that you can accomplish once you've characterized your Torsional Oscillator both mechanically and electromagnetically. Here are a few of them:

2.7.1 What is magnetic torque *in*dependent of?

In section 2.2, you found that magnetic torque would change the equilibrium position of your rotor, from an initial torque-free value to a final location where magnetic torque got balanced out by elastic torque from the torsion fiber. In that section, you found the explicit dependence of this torque on current in the coils, and later the implicit dependence of the torque on the relative angular orientation of the magnet and the coil system. But now it's time to think about what factors *don't* matter in this torque calculation.

You could ask, for example, what effect the rotational inertia of the system has -- does the final equilibrium angle attained change, if the torque is acting on a rotor with larger rotational inertia? Some things *do* change in this case, but rather than argue about whether the final angular position is one of them, you can test the matter empirically.

Similarly, you could ask what effect the damping in the system has -- does the final equilibrium angle attained change, if the torque is acting on a rotor experiencing larger damping? Here too you can do some empirical tests -- try applying the current and the resulting torque first, and then increasing the damping, and compare to the case where you increase the damping first, and only then turn on the current to get some torque. Or, you can compare the action of the magnetic dampers with hand-damping of the system -- does the choice of a damping method affect the final equilibrium position?

Finally, while you're playing with magnetic torques and damping, you owe to yourself a preview of 'critical damping'. If you're set up with not too thick a fiber, and not too much rotational inertia on the rotor, you'll find that you can adjust the magnetic dampers so as to attain critical damping -- motion in which the approach to equilibrium is *without overshoot.* (If you can attain more-than-critical damping, try adding some more rotational inertia to your system, until you get as weighty a system as you can still get to reah the critical-damping condition.) Now you're about to experience a type of motion rarely visible in the friction-dominated everyday world. Go ahead and try suddenly adding, or suddenly removing, a substantial current in your coils. You'll see the rotor swing from an initial position to a markedly different final position, but it will do so in an almost eerily smooth way -- with no overshoot, and no sudden stop at the end of its motion. Rarely will you see massive objects behave this way in your everyday experience.

2.7.2 Small oscillations about an equilibrium position

In section 2.3, you found the condition for stable equilibrium of a rotor subject to two torques -- the magnetic torque created with a current, i, and the elastic torque that develops when the fiber is twisted. You found a relationship in which the equilibrium displacement, θ, is a function of i, in which not θ, but rather $\theta/\cos\theta$, is a linear function of the current, i.

But there is more to be learned about a system than just the location of its equilibrium position -- many systems will yield *more* information, encoded in the frequency of small oscillations about a stable equilibrium position. In this project, you get to investigate such a case both theoretically and experimentally. The effect you'll see is most dramatic if you're set up with one of the thinner torsion fibers, and the smallest rotational inertia for your rotor.

To see the effect, let θ_i represent the angular displacement of stable equilibrium due to the use of current, i, in the coils. Thus θ_0 represents the equilibrium displacement of zero, when the current is off. But from section 1.3, you know you can set up oscillations about this (stable) equilibrium position, and that you can measure the period, T_0, of those oscillations. Now let θ_1 represent the equilibrium displacement you get for a current of (say) 1 Ampere in the coils. You've measured that number before, but you were perhaps deliberately uninterested in seeing the oscillations about this equilibrium value. Now you *are* interested, so learn how to set up small oscillations using by-hand intervention, and reduce magnetic damping to a minimum to make the oscillations last as long as possible. [Remarkably enough, the oscillations will last longer still if you use a constant-current, rather than a constant-voltage, source to provide the coil current i.(yes, it does make a difference! but why?)] So measure the period, T_1 or $T(\theta_1)$, of these oscillations (in the small-amplitude limit), and see if you can see a difference.

If you work at this systematically, you'll see oscillations of period $T(\theta_i)$ for small oscillations about the equilibrium position set up by current, i, and you'll find $T(\theta) < T(0)$. If you graph $[T(\theta) / T(0)]^{-2}$ as a function of $(\theta/\tan\theta)$, theory says you should expect a straight-line plot. Work out the theory, get its predictions for slope and intercept of such a plot, and compare with your data.

You'll have worked through a case of 'linear response theory' for a non-trivial mechanical system. Belonging to the same field is the diagnosis of samples of trapped atoms or trapped ions, or even whole engineering structures, by finding and interpreting the frequencies of oscillations of various modes of vibration of the system about its stable equilibrium configuration.

2.7.3 A magnetic-interaction mystery for you to solve

In this exercise, you need only the ability to make careful measurements of the oscillation frequency of your Oscillator. To make such measurements, you certainly want the oscillator to be minimally damped, but you have a choice of methods for aiming toward precision of the order of 0.1%. The simplest method is careful eyeball timing, by a stopwatch, of (say) 100 full cycles of oscillation, which might just reach the precision goal. Of course you can electronically record the position signal, and then fit it to a sinusoid, and reach higher precision still. At this level, you'll want to pay attention to energizing the torsional oscillation of the rotor, without exciting the side-to-side vibrations that might interfere with getting a good record of a sinusoid.

Once you have a demonstrated capability of precise measurement of the period, the mystery for your contemplation is simple. You just (carefully) pick up the whole Torsional Oscillator, rotate it (base, case, and all) by 180° about a vertical axis, and set it back down on the table with its feet re-occupying the same square. Now you re-measure the period, by the same protocol you've been using.

You will, in general, find a systematic difference between the two values of period. You'll find the effect on period is largest, of order 1%, if you're using the thinnest fiber, and smaller for the thicker fibers. If you can't think of why this effect exists at this level, you are free to organize your creativity by deliberately inventing some physical interaction that will give an effect of the same character. Once you discover or intuit the likely mechanism of the original mystery, you'll be able to think of ways to confirm your hypothesis.

2.7.4 Perpendicularity, and the fiber's 'angular adjuster'

In section 2.3 you modeled the effect of a current i flowing in the Helmholtz coils, creating a field, $B = k\, i$, in the magnets' vicinity, and thereby a torque, $\tau \approx \mu B = \mu k I$, on the rotor. You showed there that in fact the torque is *not* linear in the current, i, if the rotor is allowed to turn through angle, θ, in response to it. You also showed that an equilibrium deflection of the rotor by angle, θ, has the quantity $\theta/\cos\theta$, and not θ itself, nearly proportional to i.

But you may have some data showing linearity that is still imperfect, and this section takes up one reason why that happens. The model predicting that linear plot in section 2.3 assumes that when i=0 and the fiber is untorqued, the magnets' moment, μ, is exactly perpendicular to the coils' axis. Under these condition, after deflection by angle, θ, the net torque is

$$\tau = -\kappa\theta + \mu k i \sin(90^\circ - \theta)$$

$$= -\kappa\theta + \mu k i \cos\theta$$

Equilibrium occurs where this net torque vanishes, and that gives you the results of your previous model. But what happens if there's a geometrical misalignment, so that the i=0 condition put the magnets' moment, μ, not 90°, but 90° + ε, away from the coils' axis? The net torque will then be given by

$$\tau = -\kappa\theta + \mu k i \sin(90^\circ + \varepsilon - \theta)$$

$$= -\kappa\theta + \mu k i \cos(\theta - \varepsilon)$$

Equilibrium now occurs at a current, I, that is proportional, not to θ, nor even to $\theta/\cos\theta$, but to $\theta/\cos(\theta - \varepsilon)$. So guess a value for the misalignment angle, ε, and see if the use of this new model, and further guesses for ε's value, will give you a plot of $\theta/\cos(\theta-\varepsilon)$ vs. i, which turns out, is more nearly linear than your previous ε=0 assumption.

Once you can detect this effect, it's pleasant to confirm or correct it. The goal is a fine adjustment of the rotor's i=0 equilibrium location, to the let the magnets' μ lie perpendicular to the actual axis of the coils. To make that adjustment, look to the top of your torsion fiber, and see where its top clamp is mounted on a knurled disc. That disc, the 'angular adjuster', can itself be rotated by $\pm 10°$ relative to the case of the instrument. To adjust the position of this disc, you'll need to loosen (a bit) the two black-headed 8-32 socket-head screws which hold it in place -- use the 9/64" Allen tool to do so. (You might also need to reduce the tension in the fiber temporarily to make the adjustment easier.)

When you rotate the top end of the fiber by (say) 2°, you should expect the equilibrium location of the rotor's position to move by 1°. This is approximately 0.02 radian, and it's a deflection visible on the radian-protractor scale. After you make such an adjustment, in the direction you intuit, you can take another set of θ_{eq} vs. i data, and see if a new $\theta/\cos(\theta-\varepsilon)$ plot gives optimal linearity for a smaller value of misalignment, ε.

Don't be appalled if an apparently perfect geometrical orientation of the magnets relative to the coils produces a non-zero ε in your fits. There is, in practice, the possibility that the permanent magnets' μ-vectors might fail to lie exactly perpendicular to their flat faces.

As a more advanced alternative method for seeing the effect of this ε-misalignment, you can look at the velocity-transducer waveform. You know from section 1.7.5 that for oscillations of non-negligible amplitude, there are good physics reasons why this velocity waveform is not a pure sinusoid. You found there the reason that, in addition to a fundamental sine wave at the rotor's oscillation frequency f, there should be in the $V_{vel}(t)$ waveform additional Fourier components at frequencies 3f, 5f, 7f, etc.

But if there is a ε-misalignment, you should be able to show, in theory and by experiment, that $V_{vel}(t)$ will also contain *even*-harmonic components. If you have a real-time way to view the frequency content, i.e., the spectrum, of $V_{vel}(t)$, you could try rotating the fiber's angular adjuster until you maximally suppress the 2f-component of the waveform. This (you ought to be able to show) will be equivalent to dialing ε to zero.

3 Damping

Ideal oscillators are modeled as displaying exact conservation of mechanical energy, but real oscillators find ways to lose energy through various 'damping' mechanisms. In the Torsional Oscillator, there is some residual level of energy loss due to intrinsic losses in the fiber itself, and other mechanisms, but now it's time to apply some distinct modes of deliberate damping.

3.0 Applying three kinds of damping

In the Torsional Oscillator, there are provisions for experimenting with three kinds of deliberate damping.

The three kinds of damping have been chosen to display quite different dependence on velocity. The first of them you have already encountered, and is provided by the non-contact 'magnetic disc brakes' that surround the copper rotor disk. Here the damping is provided by the eddy currents that get induced in the copper, as the material moves through the static magnetic field between the jaws of the damper structures. You can adjust the degree of immersion of the copper in the field, so as to vary the damping from negligible to beyond-critical. (If you ever want to reduce this damping to a minimum, you can remove the entire magnetic-brake structures from the Oscillator's wooden box, by using the brass thumbnuts on the outer sides of the box.)

It is not at all easy to derive from first principles the amount of magnetic damping you expect in moving a conductor through a field, but it is feasible to show that the damping is expected to be very closely *linear* in the relative velocity involved. This v^1-law is also in agreement with a $v=0$ limit of no damping force at all, as is to be expected for the absence of magnetic force on non-magnetic copper. Finally, theorists love the v^1-law, not so much because it is fundamental or universal, but because it simplifies so dramatically the mathematical treatment of oscillations.

There are two other forms of damping that are easy to apply to your Oscillator: they give drag forces that vary approximately as v^0 and v^2. The velocity-independent frictional force is provided by sliding friction. The photo and diagram below show how two lines under tension, contacting the upper hub of the rotor, can be used to generate this form of damping. Again, it's not trivial to connect the coefficient of kinetic friction between the taut lines and the hub, to the numerical form of the frictional force law, but it is easy to vary systematically the tension in the lines to change the size of the damping force.

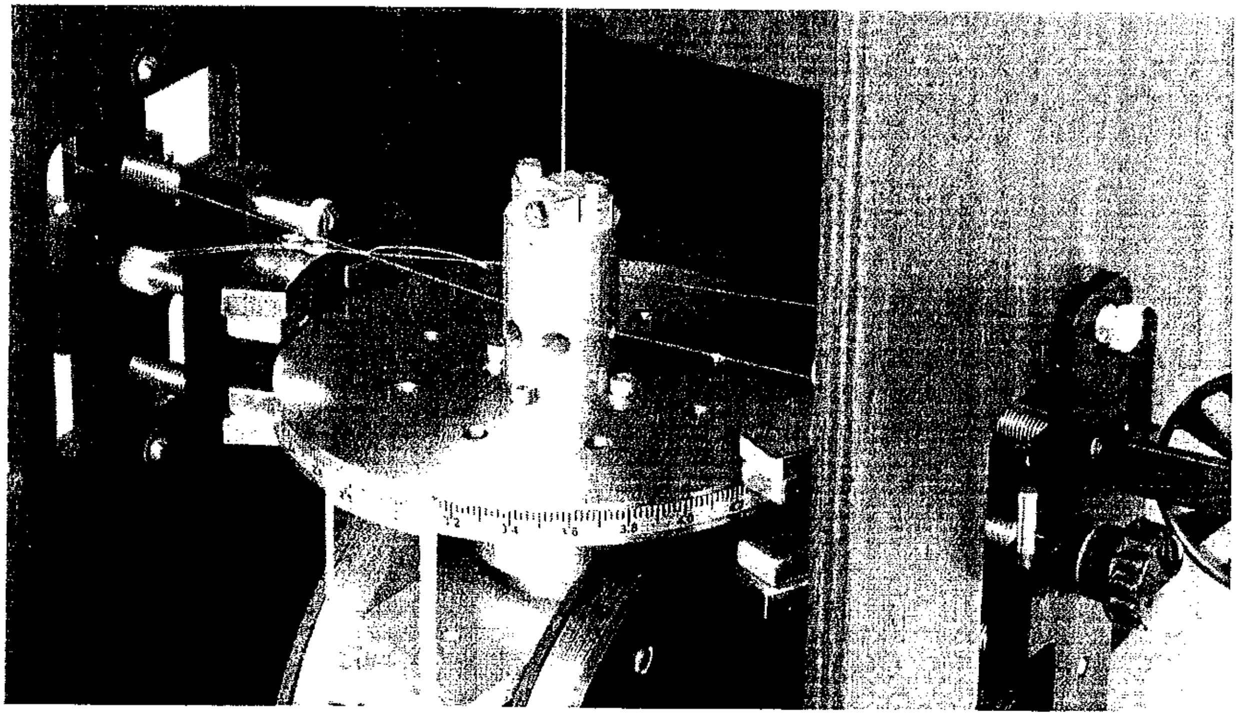

Fig. 3.0a: Taut lines in place to provide sliding friction on the rotor

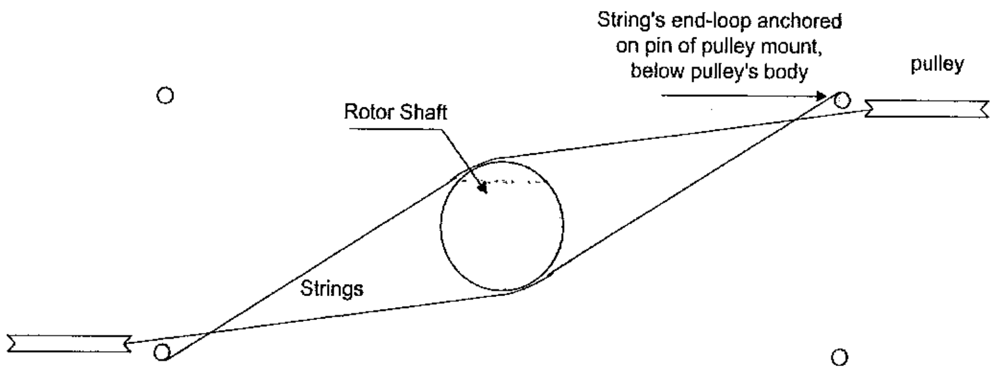

Fig. 3.0b: Schematic diagram for anchoring and tensioning taut lines

The damping that varies (approximately) as the square of the velocity is provided by fluid friction, generated when two lightweight 'paddles' are attached to the hub of the rotor. Particularly for large amplitudes of oscillatory motion, the speed of the paddles through the air puts the fluid friction in a regime where v^2-dependence is expected. Then the usual empirical model involving a 'drag coefficient' should give a decent representation of the frictional force. Here too the drag force can be varied, since the paddles can be 'feathered' (i.e., rotated by 90° about the axis of their long tubes) to reduce the frontal area, and thus the frictional force, dramatically.

It's worth mentioning that the paddles have been built as light as possible, but even so you'll find their addition will markedly raise the rotational inertia of the rotor. Because the aluminum tubing that forms their arms has a very thin wall, you should **not** over-tighten the nylon thumbscrews that hold these shafts into the hub of the rotor -- just barely finger-tight is right in this application. Similarly the foam-core board that forms their vanes could be easily crushed by contact with something stiffer than air. The paddles can be safely stored using the eye-hooks and L-hooks on the back of the Oscillator's box when they're not in use.

3.1 Linear damping

You've seen the 'magnetic dampers' in action in previous sections, and now it's time to describe and discover what the presence of a linear-in-velocity law does to oscillatory motion that would be otherwise purely periodic. In doing these investigations, it's crucial to have a digital oscilloscope, or other data-recording tool, to acquire records of the angular-position transducer as a function of time. Since damping will only remove energy from the oscillator, you'll need other ways to put energy in. This can be either hand-excitation of the copper rotor disc itself or of the little 'pumper-upper' wire clamp mounted low on the fiber, or it can be from currents injected into the Helmholtz coil as a torque drive.

The Oscillator will display its simplest possible behavior if it's held away from equilibrium, and then released from rest. The cleanest 'hold and release' technique is to put some steady current into the drive coils, let the rotor settle down at some displaced equilibrium position, and then reduce the current suddenly to zero. Labeling the time when the current drops to zero as the $t=0$ point, you should see the position signal possessing some steady non-zero value for $t<0$, and then exhibit 'damped oscillations' for $t>0$.

It is a characteristic of a v^1 damping law that the rate of energy removal from the system is proportional to the energy already present. Such a relationship ensures that the mechanical energy in the system decays exponentially. So the oscillations in time of a position waveform ought to fall within an 'envelope' which itself displays exponential decay. You can follow those oscillations from radian-scale to milli-radian size. In the next section you'll learn how to model mathematically the data you acquire.

The rate-constant of exponential decay varies as you change the degree of magnetic damping via the adjusting knobs on the dampers. You might want to position the two dampers so that they engage the copper to equal degrees. Make fine adjustments, or repeatable settings, of the dampers' positions, by counting full (and partial) turns of the knobs. Each full turn of the knob will move the damping structure by 1/20 of an inch, i.e., 1.27 mm. If you mark the knob and think of a clock-face, estimating partial turns to the 'nearest hour' will give you control at the 0.1-mm level.

By withdrawing the dampers maximally, or even removing them entirely, you can check what 'baseline' damping comes from other loss mechanisms. That intrinsic damping is not so simple as the magnetic damping you've been studying. It is apparently dominated at higher amplitudes by friction of the copper disc's motion through the air, and at low amplitude by losses inside the metal fiber. Notice that under these conditions of minimum-possible damping, the envelope of the oscillations might not be a single simple exponential.

3.2 Modeling damped oscillations

In a regime of moderate magnetic damping, your data can be modeled by an exponential damping law. Section 3.1 has had you take the data, and this data can now be compared to theoretical expectations.

You have previously measured the period, T, of oscillations, and for damping linear in the velocity (called 'linear damping', even though it yields an exponential decay law) the zero-crossings of the position signal are still expected to be strictly periodic, with some period, T. But this period is predicted to be subtly different from the period T_0 that would be expected for an undamped oscillator. The theory is traditionally worked out in theorists' notation, which assigns 'free angular frequency', $\omega_0 = 2\pi/T_0$, to the undamped motion, and 'damped angular frequency', $\omega_d = 2\pi/T_d$, to the damped motion. [Here the ω-values have units of radians per second, or just s^{-1}, for which the notation Hz is *not* used. The ordinary frequencies are not measured in radians (of phase, in the phase plane) per second, but in full cycles of oscillations per second, so $f_0 = 1/T_0$ and $f_d = 1/T_d$ give the frequencies that are properly reported in Hertz.]

Now the expected form for damped oscillatory motion can be modeled mathematically in terms of an 'undamped frequency', ω_0, and a (dimensionless) 'damping coefficient', γ. The combination of exponentials and sinusoids that is expected to appear is

$$\theta(t) \propto (\text{const}) \exp(-\gamma\, \omega_0\, t) \cos[\omega_0\, t \sqrt{(1-\gamma^2)}]$$

Thus the damped frequency is predicted to be slightly *lower* than the undamped one, by a factor of

$$\omega_d = \omega_0 \sqrt{(1-\gamma^2)}$$

The experimenter's approach to a set of data is to try fitting the data, in the least-squares sense, to a function of the form

$$\theta(t) \approx A \exp(-B\, t) \cos[C\, t]$$

or better,

$$\theta(t) = A \exp(-B\, t) \cos[C\, t - D] + E$$

In such fits, the coefficients A (and D) will depend wholly on initial conditions, but the B- and C-constants are predicted to be *in*dependent of initial conditions, and also to be related according to

$$B^2 + C^2 = \omega_0^{\,2}$$

So while the fitting parameters B and C will vary with the degree of damping you choose, the sum of their squares is predicted to be independent of damping. To get a visual indication of this, you can get combinations of B- and C-parameters for data sets with different degrees of damping, and then plot pairs (C,B) in a Cartesian plane -- theory predicts all the points will fall on a quarter-circle.

If you try taking this data, you'll find it takes a very high degree of damping to get points that 'fill in' the quarter-circle of points. In fact you'll close in on 'critical damping', defined to lie at $\gamma = 1$ in this theoretical description. You can show that critical damping is the degree of damping that's just sufficient to suppress any overshoot in the approach to equilibrium, or alternatively the damping required to ensure approach to equilibrium from one side only, without oscillations.

Fits like these, to data obtained with less extreme damping, will enable you to *falsify* the assertions of those authors who claim (noting that ω_d and ω_0 are very close together for modest damping, and that oscillations die away very quickly for larger damping) that ω_d and ω_0 are indistinguishable in practice.

With a fitting technology in hand that will produce accurate values of

$$B^2 + C^2 = \omega_0^2$$

you are in position to test the repeatability and the precision with which you can measure the 'natural frequency' ω_0 of your oscillator. Try some successive trials under nominally identical conditions to see how reproducible your values of ω_0 are. For a reality check that you are measuring something that could vary, try placing a little washer, or a paper clip, atop the edge of your rotor disc -- this should raise its rotational inertia (by how much?), and thereby raise the period, and lower the natural frequency, of your oscillator. Given the observed repeatability of measured ω_0-values as a 'noise level', how small a 'signal', in the form of a change in natural frequency, could you reliably detect? There are lots of sensors in the world which depend, for their sensitivity, on the ability to detect a small change in the natural frequency of some oscillating system.

3.3 The 'Q' of a damped system

So far you've modeled damping by a dimensionless parameter, γ, defined so as to describe undamped motion for $\gamma = 0$, critically damped motion for $\gamma = 1$, and damped oscillatory motion for $0 < \gamma < 1$. (The case of $\gamma > 1$, over-critical damping, is addressed in section 3.6). This section introduces you to a very common way of describing oscillatory motion of low damping -- that is via the 'Q', or Q-factor, or 'quality factor', of an oscillatory system.

One formal definition of the Q of a system is via

$$Q \equiv 1/(2\gamma)$$

This assigns the value Q = 0.5 to a critically-damped system, and has the Q rise to 1 for a less-damped, and Q > 1 for an even-less-damped system. In fact the Q-factor rises rapidly as the damping gets very small, and the Q is typically given as a figure of merit for systems of very low damping. Different branches of physics boast systems of various Q-values, with record Q's of 10^6, 10^9, 10^{12} or even higher. These stand for systems of amazingly low damping, ever closer to the 'Platonic ideal' of an undamped oscillator.

The Q reappears in various guises, in Section 4 of this manual, in connection with various resonant responses of a damped oscillator. For now, work with the mathematical model of damped motion in section 3.2 until you can connect the parameter, γ, the Q, and $N_{1/2}$, the number of cycles of oscillation that go by while the oscillations are decaying to half their original amplitude. You should be able to show that

$$Q \approx 4.53\ N_{1/2}$$

which gives a great way to estimate the Q of any oscillating system whose decaying waveform can be followed through a factor-of-2 decay in amplitude.

Since high Q's earn you bragging rights in certain contexts, see how high a Q you can get in your Torsional Oscillator. Clearly magnetic damping can be reduced to near-zero, leaving air-drag and fiber-loss damping only. Certainly any slippage in the clamps that hold the fiber will worsen your Q. You may also find that no single Q-value describes the full motion that remains under these circumstances. The highest Q occurs in oscillations of quite small amplitude. Nor is it clear if smaller or larger diameters of fiber, or greater or lesser values of rotational inertia, will give the highest possible Q. There have certainly been very detailed studies of the highest Q's achievable in various torsional oscillators, and extreme cases use sapphire fibers, with special surface treatment, at cryogenic temperatures, in vacuum, to achieve the highest possible Q-values. Steel music wire, at room temperature, in air, won't come close to competing in this league, but you may still be able to demonstrate a Q of 1000 or more on your tabletop apparatus.

3.4 Finding critical damping

This section deals with the phenomenon of critical damping, easily achieved with the use of the magnetic dampers on the Torsional Oscillator. It expands on the modeling of section 3.2.

The model of damping that displays exponentially-damped sinusoids can be written as

$$\theta(t) \approx (\text{const}) \exp(-\gamma\, \omega_0\, t) \cos[\omega_0\, t\, \sqrt{(1-\gamma^2)}]$$

This model depends on a choice of t=0 that is not guaranteed to match the 'initial condition' of your oscillator. The simplest 'hold and release' experiments can be described, in the language of differential equations, as a case in which

$$\theta(t=0) = A$$

but

$$d\theta/dt(t=0) = 0$$

That is to say, the oscillator is released from a non-zero position, A, but with zero initial velocity. Under these conditions, the solution to the linearly-damped oscillator problem is given by

$$\theta(t) = A \exp(-\gamma\, \omega_0\, t) \cos[\omega_0\, t\, \sqrt{(1-\gamma^2)} - \delta] / \sqrt{(1-\gamma^2)}$$

where δ is an angle, or phase shift, in radians given by

$$\tan \delta = \gamma / \sqrt{(1-\gamma^2)}$$

This is equivalent to

$$\cos \delta = \sqrt{(1-\gamma^2)}$$

so another form of this result is

$$\theta(t) = A \exp(-\gamma\, \omega_0\, t) \cos[\omega_0\, t\, \sqrt{(1-\gamma^2)} - \delta] / \cos \delta$$

It's a purely mathematical task to show that this form of the solution matches the initial conditions specified.

One physical use of this result, which now assigns the t=0 location as the instant of release, is to look physically, and mathematically, for the time of the first zero crossing, which is set by requiring

$$\cos[\omega_0 t \sqrt{(1-\gamma^2)} - \delta] = 0$$

or

$$\omega_0 t \sqrt{(1-\gamma^2)} - \delta = \pi/2$$

This assigns a value to the time of the first zero-crossing, $t_{\#1}$, which is very close to a quarter-period of the undamped oscillation in the limit of small damping, but which grows quite markedly as the damping gets larger. In fact, as the damping approaches critical, the value of $t_{\#1}$ goes out to infinity, which is to say, the first zero crossing never occurs at all.

A series of 'hold-and-release' data sets, taken at increasing levels of damping, will show this 'divergence' of $t_{\#1}$. To learn where just-critical damping occurs, tabulate instead the *inverses* of these numbers and plot $1/t_{\#1}$ values, because these values ought to extrapolate to zero at the attainment of critical damping.

Finally, a great deal is made, mathematically, of separate treatments of sub-critical, exactly-critical, and over-critical damping, which might lead you to think there is some discontinuity in the *physics* of the situation. That's not true in practice, since you can make damping adjustments in a perfectly continuous way from one side of criticality to the other. What you're seeing is only a 'discontinuity' in the mathematical *language* that's convenient for modeling the motion. And while you'll never be able to dial the damping parameter γ exactly to 1, you will still find it a convenient target to aim for, particularly when in section 3.6 you use it to attain a known degree of *over*-criticality in damping.

3.5 Step response and impulse response

This section introduces two very widely used vocabulary terms for damped oscillators in general. Thus far you've been seeing the behavior under 'hold and release' conditions, but there are much more general conditions that can excite the oscillator. These experiments require an external current source to energize the 'torque drive' or Helmholtz coil, one that one can be turned on and back off for a known and variable duration.

The first sort of drive waveform is called a 'step function', in which the drive current changes from $i=0$ at times $t < 0$, to a constant $i = i_0$ for times $t > 0$. The system starts with zero displacement, and zero velocity, right up until $t=0$. To keep the response of the oscillator in the small-angle regime, motivated in turn by the complications in the torque drive that you encountered in section 2.3, use currents i_0 under 1 A.

In response to this one-time discrete 'step' in the current, or the drive, or the 'cause', you'll see an 'effect' or a response which is not as simple as a single step function. In the case of damping below critical, you'll see overshoot, oscillation, and eventual asymptotic approach to a new equilibrium position. For dimensionless damping parameter γ lying in the range $0 < \gamma < 1$, the predicted form of the response is

$$\theta(t) = \theta_\infty \{ 1 - \exp(-\gamma\, \omega_0\, t) \cos[\omega_0\, t \sqrt{(1-\gamma^2)} - \delta] / \cos \delta \}$$

where θ_∞ is the long-term limit of the angular response, and where δ is an angle given by

$$\tan \delta = \gamma / \sqrt{(1-\gamma^2)}$$

In one sense, there's nothing new to this shape of response compared to your former hold-and-release data. Here you are instead 'holding' the system at $\theta(t=0) = 0$, and then 'releasing' it to seek the equilibrium it'll have in response to a drive current i_0. Extract, by fitting, values of the parameters γ and ω_0 from the step-function response, just as you did formerly in case of hold-and-release.

More interesting, perhaps, is the impulse response, which requires a different sort of excitation. Start with a quiescent oscillator, and turn on the current to a value of i_0 at $t=0$. But now instead of holding it indefinitely at that value, return the current to zero after a finite, and rather short, time, t_f. Finally, imagine that you make t_f smaller and smaller, but you raise the value of i_0 in proportion to this shrinkage, to create a set of excitation pulses that more and more nearly approximate a 'delta function' of negligible duration but fixed area. In actual experimental practice, you don't need to reduce t_f to zero, but you do need to achieve

$$t_f/T << 1/2\pi$$

where T is the period of oscillation of the system.

So, generate an approximation to an impulse function, and then try using double the current for half the duration to see that the response of the system is very nearly identical. What you're seeing is called the 'impulse response' of your oscillator, and its waveform in time is predicted to be proportional to the *derivative* of the step-response waveform. One way to write the impulse-response function analytically, for the case of an oscillator that's less than critically damped, is

$$\theta(t) \propto \exp(-\gamma\ \omega_0\ t)\ \sin[\omega_0\ t\ \sqrt{(1-\gamma^2)}]\ /\ \sqrt{(1-\gamma^2)}\ ,$$

and once again you can fit a mathematical function to the observed impulse response so as to extract the parameters γ and ω_0 for a system.

The special borderline case of a critically-damped system gives a simpler-still impulse response, since the function above can be shown to have the $\gamma \rightarrow 1$ limiting behavior of

$$\theta(t) \propto t\ \exp(-\ \omega_0\ t)\ \ [\text{for the case } \gamma = 1].$$

This form makes it clear that the impulse response starts at zero, departs from zero, and returns asymptotically and exponentially to zero, with no other zero crossings in the interval $0 < t < \infty$.

The remarkable thing about step- and impulse-responses is that either waveform, observed experimentally *for any linear system*, gives in some sense all the information that is required to model the system fully. Another way to say that is, *for any linear system*, knowing just the unit-step response suffices to enable the prediction of the response of the system *to any form of excitation whatever*. Thus the experimental measurement of the step response of an experimental system is a diagnostic of very wide generality and usefulness.

3.6 Over-critical damping and how to achieve it

You've learned in section 3.2 how to model a linearly-damped oscillator that's below critical damping and in 3.4 how to find the point of critical damping. This section, still using the v^1-law magnetic dampers, suggests a method for investigating over-critical damping quantitatively. In performing these exercises, it's desirable be using one of the thinner of the torsion fibers, and to start with a rotor bearing maximal rotational inertia.

The idea depends on changing the rotational inertia of the system, to change from a system that's just critically damped, to one that's over-critically damped, and by a known amount too. Suppose that for an oscillator of torsion constant, κ, and rotational inertia, I, the dampers have been adjusted to give a damping torque

$$\tau = - b\, d\theta/dt$$

opposite to and proportional to the angular velocity, with a *b*-coefficient that is just sufficient to make the damping critical. Since the dimensionless damping parameter, γ, is predicted to be

$$\gamma = b / [2 \sqrt{(\kappa I)}\,]$$

this procedure will have empirically set $\gamma = 1$, or will have adjusted the dampers until

$$b = 2 \sqrt{(\kappa I)}$$

Now consider taking masses <u>off</u> of the rotor. That won't have changed κ (which depends only on the choice of fiber), and it won't have changed *b* (which depends only on the overlap of dampers and copper rotor disc), but it <u>will</u> have changed I, to a new and smaller value, I'. In fact, via the modeling of section 1.3, you can even say numerically what the values of I and I' are. Now the new damping parameter will be

$$\gamma' = b / [2 \sqrt{(\kappa I')}\,] = b / [2 \sqrt{(\kappa I)}\,] * \sqrt{(I/I')} = 1 * \sqrt{(I/I')}$$

So if removing some masses were to have halved the rotational inertia, so

$$I' = (1/2)*I$$

then the system would have changed from $\gamma = 1$ to

$$\gamma' = 1*\sqrt{(I/0.5*I)} = \sqrt{(2)}$$

So it will be *over*-critically damped, and by this computed amount. The system will also have a new value of 'natural frequency', given by

$$\omega_0' = \sqrt{(\kappa/I')} = \sqrt{(\kappa/I)} * \sqrt{(I/I')} = \omega_0 * \sqrt{(I/I')}$$

The payoff of this change is that it leads to mathematical predictions for the behavior of an over-damped system. Instead of getting a solution described mathematically by the product of a decaying exponential and a sinusoid, the claim is that the solution will be described as a sum of two decaying exponentials -- one decaying faster, the other decaying more slowly. In fact, the exponential decay rates are predicted in detail, giving the claim

$$\theta(t) = c_1 * \exp\{-\omega_0' [\gamma' + \sqrt{(\gamma'^2-1)}]\, t\} + c_2 * \exp\{-\omega_0' [\gamma' - \sqrt{(\gamma'^2-1)}]\, t\}$$

where c_1 and c_2 are constants determined by the initial conditions.

So using previously measured values of ω_0 and I, and the new (and smaller) value of I', you should have gone from critical damping, $\gamma = 1$, to new and computable values of ω_0' and γ'. Acquire data in the over-critically damped condition, and model it by a $\theta(t)$ equation of the form above -- that is to say, see if the predicted decay rates, and correctly chosen (or fit) values of the constants, c_1 and c_2, will describe the data that you get. For quite modest degrees of over-critical damping, observe that the two exponential decay rates are dramatically different. Notice too that the more quickly-decaying exponential becomes negligibly small, compared to the more slowly-decaying one, after a surprisingly short time.

3.7 Sliding friction

Finally it's time to change from the v^1-law magnetic damping to the v^0-law expected in the case of sliding friction. In section 3.0 we described one way to get a controllable and modest amount of sliding friction in your oscillator, and this section describes some observations you can make under these conditions.

The mathematical model of sliding friction is more complicated than it would first seem. It is *not* correct to assume a force given simply (say) by $F = - b\, v^0$, or a torque given by

$$\tau = - b\, [d\theta/dt]^0$$

since though sliding friction might have a magnitude independent of the size of the velocity, it always conspires to have a *sign* that is opposite the instantaneous velocity. In the case of sliding friction in one-dimensional motion, you can model this mathematically by assuming

$$F = - b\, v/|v|$$

which accomplishes the description desired.

More fundamentally, departing from a v^1-law removes the mathematical property of *linearity*, which means that in the case of sliding friction, there is no one general solution. In particular, you'll find that initial conditions matter more than in just setting a scale factor.

The data simplest to collect are from hold-and-release experiments, conducted with the magnetic dampers fully withdrawn (or entirely removed), but with two taut lines in position to give sliding friction. Whether by mere viewing, or by electronic data-acquisition, you'll rapidly see that the angular position coordinate, θ (t), undergoes decaying oscillations of a character quite different than you've seen so far. Rather than oscillating indefinitely (albeit at ever-lower amplitude), these new motions are of finite duration, and come to a definitive stop. (And why do they stop, suddenly? And do they stop at $\theta = 0$, or somewhere else? And what friction law is acting <u>after</u> they stop?)

You can vary the amount of sliding friction in two ways. The simpler is to change the tension in the taut lines, by adding more mass to the hangers that are making the lines taut. The other way is to increase the length of the arc over which the string is in contact with the rotor's hub, perhaps by changing the anchor point of the fixed end of the string. Either method will make a difference in the motions, but either way, you'll still find a qualitative difference between these oscillations and the ones characteristic of linear damping. Pay attention to the *envelope* of the oscillations, and remember this as a signature of damping by sliding friction, or indeed any damping force which does not vanish in the $v \to 0$ limit.

There's more modeling you can do, beyond these important but rather qualitative distinctions you've seen, but that's described in section 3.10.1.

3.8 Fluid friction

Section 3.0 described how to apply various kinds of damping to the otherwise nearly-undamped motion of the Torsional Oscillator. This section takes up the case of damping from fluid friction, obtained using the lightweight 'air paddles' that can be mounted to the hub of the rotor. In performing these experiments, you'll want to remove any sliding friction due to taut lines, and you might want to withdraw or remove the magnetic-damper assemblies too.

Hand-exciting some rotational motion of the Oscillator with the dampers in place, two novelties should be apparent. First, there's a reduced amplitude range to investigate, since the arm of the rear paddle can only swing through so long an arc before colliding with the wooden case of the instrument. Second, the very modest masses of the paddles have added a surprisingly large amount to the rotational inertia of the system, and the 'natural frequency' has accordingly dropped considerably. You can use the methods of section 1.3 in reverse to measure that increment in rotational inertia, or the data on the paddles in section 6.1 to predict it.

Now the character of the motions with the paddles in place depends a great deal on the amplitude of the oscillations. At small amplitudes, the paddles seem to do little (except for that notable change in rotational inertia). To check this, rotate the paddles by 90° about the long axes of the aluminum shafts, 'feathering' their surfaces, and see how little difference that change makes.

By contrast, you can perform hold-and-release experiments starting with the *maximum* amplitude that your geometry will allow. If you record the time evolution of the angular coordinate $\theta(t)$, you'll now find that 'feathering' now makes a huge difference. So now the paddles, working face-on, are creating a large effect on the motion of the system. It is also a motion that is hard to describe in any generality -- specific detailed modeling is deferred to section 3.10.2. But look at the *envelope* of the oscillations, and compare it to the envelopes for the earlier cases of v^1-law and v^0-law damping. Fluid friction gives very fast damping of the early (i.e., large-amplitude, high-velocity) motion, and very little damping of the late (small-amplitude, low-velocity) motion, suggesting that the damping's dependence on velocity is of higher order in the velocity. The simplest models of air resistance suggest a v^2-law, consistent with these qualitative observations.

Once again, damping of this form removes the mathematical property of linearity from the system. In particular, there is no longer any guarantee that the oscillations are even 'isochronous', and in particular, you might expect zeroes of the angular position during early half-cycles of the motion to have a *different spacing in time* than the values you get later in the motion. This effect on the period of the system is another feature that detailed models of section 3.10.2 ought to reproduce.

3.9 Energy modeling of damped motion

You've now seen examples of damped motion that are more or less well-described as arising from frictional-force laws of v^0-, v^1-, and v^2-character. This section aims at a unified description of friction's effects, according to an assumed v^n power-law, to give you a new way to diagnose damped motion, even when the mechanism of damping is unknown.

Suppose there's present some kind of damping, but a case in which the damping is not all that strong. Suppose, in fact, that the qualitative description of the motion is still oscillatory, with a succession of half-cycles of the motion, only with decreasing amplitude. The idea is to compute the energy of the system, and isolate the loss of energy that can be blamed on damping.

It's worth thinking about half-cycles of the motion, from a (positive) maximum value of $\theta_1 = \theta(t_1)$, to a (negative) minimum value of $\theta_2 = \theta(t_2)$. That's because in the time interval $t_1 < t < t_2$, the velocity is always of one (negative) sign, and that means the assumed damping torque must have been positive during that time interval. According to our assumptions, we can write it as

$$\tau = (+)\, b\, |\, d\theta/dt\, |^{\,n}$$

where b is some positive constant, and the correct sign has been put in by hand.

Now for an energy accounting. For any motion that's of oscillatory character, there are a succession of times (t_1, t_2, t_3, t_4, . . .) at which the angular coordinate, θ, reaches extrema, alternately positive and negative. At each such extremum, by definition, the kinetic energy of the system is zero, and so all of the energy is of elastic-potential in character. That energy value is given by

$$E = (1/2)\, \kappa\, \theta^2$$

where the torsion constant κ has a known value. So a set of data can be used to quantify values of the system's energy at these times, giving E_1, E_2, E_3, E_4, . . . In each half-cycle,

$$\langle E \rangle = (E_n + E_{n+1})/2$$

is a fair approximation of the average energy of the system during the entire half-cycle, while

$$\Delta E = |E_n - E_{n+1}|$$

is the amount of energy that has disappeared from the system during that same half cycle. Now a pair of numbers, (<E>, ΔE) has been associated with each half-cycle of the motion, and the new idea is to plot the pairs (<E>, ΔE) for lots of half-cycles in a Cartesian plane. The points thereby plotted do *not* lie at random in the plane, but fall along a characteristic curve. And the curve along which they lie will *differ* in the case of the three different damping mechanisms that you've investigated experimentally.

The theory of those curves requires a bit more effort. That starts with the form of the work-energy theorem for rotational motion, which can be written as an integral, over the chosen time interval, of the product of torque and angular displacement:

$$W = \int dW = \int \tau\, d\theta = \int \tau\, (d\theta/dt)\, dt = \int b\, |d\theta/dt|^n\, (d\theta/dt)\, dt = (\pm)\, b \int |d\theta/dt|^{n+1}\, dt$$

Here the result will give the work done by the damping force in half a cycle, if the integral is taken over half a period. And in the weak-damping limit, one can assume that a half-cycle of motion is described by

$$\theta(t) = A \cos \omega(t - t_1)$$

$$d\theta/dt = A\,(-\omega) \sin \omega(t - t_1)$$

where A is the current size of the amplitude, and ω is the natural frequency of the current motion, with the current half cycle starting at time t_1. With this 'current cycle' approximation, the integral can then be performed for various n-values to give the energy loss per half-cycle. The results can be used to show why the energy of the system decays exponentially (for the case n=1, v^1-law damping appropriate to the use of magnetic dampers), or why the *amplitude* of the system decays linearly (for the case n=0, v^0-law damping appropriate to sliding friction). What results emerge for n=2, a v^2-law of damping? And how well does this model fit your data for fluid-friction damping, plotted via those ($\langle E \rangle$, ΔE) points in a plane?

This energy accounting is not the full story, but since it gives some way to understand the time evolution of the system's energy with*out* the need to solve in detail (and numerically) the differential equation describing it, it is an valuable approximate description of damping (provided it's not too strong).

3.10 Projects

With all the variety of techniques you now know, there are lots of projects involving damping that you can accomplish.

3.10.1 Mathematical modeling of sliding friction

This section suggests two ways to understand the effects of sliding friction, modeling it as giving a torque whose magnitude is velocity-independent, but whose sign is always opposite the present velocity.

Both methods require setting up and solving the differential equation describing the motion. Using Newton's Second Law for rotational motion in the form

$$I\, d^2\theta/dt^2 = \sum \tau = -\kappa\,\theta + \tau_{damping}\,,$$

you need only put in a form for the damping torque to get a complete equation.

A form best suited for numerical solutions is to write

$$\tau_{damping} = -b\,(d\theta/dt)\,/\,|\,d\theta/dt\,|\,.$$

where the $z/|z|$ form correctly gives the sign and the magnitude of the damping torque. Of course an actual numerical solution requires numerical values of I (the rotational inertia of the rotor), κ (the torsion constant of the fiber), and b (the coefficient of the damping term). Only the last of these has to be guessed, and of course the numerical method of choice can be validated by testing it first using b=0. Solving differential equations numerically also requires assuming initial conditions, but values for the angular position and angular velocity appropriate to hold-and-release conditions are easily found. The only thing that can go amiss with the numerical solution is the indeterminate form that arises at any instant at which the angular velocity happens to be zero. Finally, the numerical result is totally particular to the numerical parameters used, and every case is another special case.

The alternative method of solution which happens to work in this case of velocity-independent damping is wholly analytical, and it handles the reversal of direction of the damping force by treating each half-cycle of the motion separately. Again, for hold-and-release conditions, we can assume (say) a positive value of angular position, and a zero value for angular velocity, and then we know (on physical grounds) that the angular velocity will be *negative* during the first half-cycle. So the damping torque has to be taken to be a constant (of size b) and *positive* during this half-cycle. That gives the differential equation

$$I\, d^2\theta/dt^2 = \sum \tau = -\kappa\,\theta + b$$

which can be solved analytically. It's an inhomogeneous equation, but it's simple to find the 'general solution' and the 'particular solution', and to make their sum fit the initial conditions. The motion is predicted to start with zero velocity, and to attain zero velocity again after a finite time. That's the time, physically, at which the motion will reverse direction, and the end of this half-cycle provides the initial conditions for the next half-cycle. During the second half-cycle, the damping torque has to be taken as the constant (-*b*), of course, to keep it opposite to the now-positive sign of the angular velocity.

Piecing together the effects of two half-cycles, show that the motion during a whole cycle returns the system to the zero-velocity condition with an amplitude *reduced* from the original release point, and that the *same* reduction in amplitude will also occur in every full cycle into the future.

What stops the motion in the end? Again, it takes some physics assumptions. If the arc of contact of taut line and rotor hub provides a maximal force of *static* friction that is larger than the value of *sliding* friction, then eventually there occurs an instant of zero velocity at which the angular displacement is small enough that the fiber's torque is insufficient to break the rotor free from the force of static friction. So static it should stay!

3.10.2 From a coefficient of friction to an energy-loss model

The effect of sliding friction has been handled, so far, with a wholly empirical coefficient, '*b*', and this coefficient is left unrelated to the more familiar 'coefficient of kinetic friction', μ_k, beloved by textbook authors. What's the connection?

To work this out, and ultimately to make your hold-and-release data correspond to a measurement of the coefficient of friction, requires a few steps. Those steps are best taken under the assumption that the angular velocity is temporarily of some single sign (say, positive), in which case you should be able to intuit that the taut line, rubbing against the hub of the rotor, is in fact not all at one single tension. Instead, the tension must be different in the two segments of the line, one of them 'upstream' and the other 'downstream' from the rotating hub.

To understand the connection between the tension in these two segments, you should look into 'Euler's capstan equation', a very elegant connection between the coefficient of kinetic friction assumed to exist in the interaction of the line and the hub, and the difference in the tensions in these two segments. (It also marked, historically, one of the first physics applications of the exponential function.) It working out the model, you will need to know the measure of the arc over which the line and hub are in contact, and this can be estimated from the angle through which the line is 'deflected' in going around the hub. For the small arc of contact you will typically use, you might find it useful to make a simple series expansion of the exponential you'll get from Euler's equation.

With those two values for the tension, and a known radius for the hub, you should see that there are two torques, of opposite direction and slightly different magnitude, acting on the hub due to the (two segments of) line. Their difference is the torque due to sliding friction, and if you've completed the above accounting, you'll have a *b*-coefficient that's been connected to an assumed coefficient, as well as other measurable parameters.

3.10.3 Modeling the v^2-law

A physicist's usual encounter with a v^2-law for fluid friction typically comes with an equation of the form

$$F_{drag} = (1/2)\, C_D\, \rho\, A\, v^2$$

which claims to give the drag force for an object of frontal (or cross-sectional) area, A, moving at speed, v, through a fluid of density, ρ. Here C_D is a (dimensionless) drag coefficient, which is presumed to depend on the shape of the object moving through the fluid. You can even find references that claim $C_D \cong 1.2$ for a 'flat plate' moving through a fluid.

But where does this equation come from? And why is it reasonable? (Note that the 'law' above ascribes the friction *not* to the viscosity, but only to the density, of the fluid.) Even these questions set to the side the hardest question of all, which is to ask -- in what velocity regime can it be expected to apply? Here's some simple analysis that suggests why, in the flat-plate limiting case, there ought to be an equation of this form at all.

Consider the plate, then, as an area, A, moving at speed, v, through a fluid, and consider a time interval of duration Δt. In that time interval, the plate moves forward a distance of v Δt, perpendicular to its face. Clearly the fluid that was formerly in the volume, v Δt, must have moved out of the way. To a first approximation, we can assume that it's been shunted aside at negligible velocity -- this pre-supposes that a pressure wave moves through the fluid well in advance of the flat plate itself, and that there's 'plenty of time' for the fluid in front of the plate to get pushed out of the way without having to acquire much velocity. (Clearly, this requires the assumption that v is far below the speed of sound in the fluid.)

Things are different with the space *behind* the moving plate, also of volume, v Δt, that has freshly opened up with the plate's movement. That space is presumably not 'full of vacuum', but is also full of fluid of density, ρ, so that there's a mass of fluid, in fact of mass

$$\rho\, V = \rho\, v\, \Delta t$$

filling that space. But for the fluid in there to be 'keeping up' with the moving plate, the fluid immediately behind the plate must itself be moving forward at speed, v, just as the rear surface of the plate is. So, goes this analysis, there is an *energy cost* to move the plate forward, corresponding to giving kinetic energy to the fluid that is continually filling in the space behind the moving plate.

If you work out the energy that this costs per unit time, and then per unit distance, you can get a 'work per unit distance' it takes to move the plate through the fluid. You can even see where that work is going -- it ends up in the extra kinetic energy of the (presumably, soon turbulent) fluid in the wake of the moving plate. And a 'work per unit distance' is dimensionally, and actually, a force. Working out the algebra of all these assumptions will indeed give a drag-force equation of the usually-advertised form, and will even give a 'proof' that the drag coefficient should be $C_D = 1$. (Improving on this prediction would be really hard computational work!)

Look up the 'Reynolds number' to understand, at least at the empirical level, something about the velocity regime in which this law is observed to be a good approximation. Then you can work out the Reynolds number appropriate to the actual conditions of your paddles, moving at their peak velocities, in your Torsional Oscillator experiments. Finally, check to see if your data from oscillations damped by the air paddles is at all consistent with a drag coefficient of order 1.

3.10.4 Oscillations damped by a v^2-law

Suppose we take a v^2-law of fluid friction as a given, and assume some empirical value of a drag coefficient. What should that do to oscillatory motion? The answer will require you to solve a differential equation numerically, since there's no real progress possible analytically.

First, the real model you want to write for a v^2-force of drag in a fluid is

$$F = - b\, v\, |v|$$

where the constant, b, is related to the drag coefficient and other parameters. The motivation for the absolute-value sign is that this gives a force whose direction is always opposite to that of the velocity, but whose magnitude has the assumed v^2-scaling.

The next task is to go from a force-on-one-paddle to a torque-on-the-rotor model, which will give something like

$$\tau = - b' (d\theta/dt)\, |d\theta/dt|$$

where b' is another constant with a numerical value. Now write the 'equation of motion' for the angular position as

$$I\, d^2\theta/dt^2 = \Sigma\, \tau = - \kappa\, \theta - b' (d\theta/dt)\, |d\theta/dt|$$

which is a well-defined, though non-linear, differential equation. Solve it numerically, subject to initial conditions that correspond to your hold-and-release experiments.

There are plenty of things to look out for. One valuable check of the numerical method is to treat the b'=0 case numerically -- since in this case of no damping, analytic methods can also give the solution. Another check is to look at late times, when the effect of the damping is small, to see if the behavior of slowly-decreasing amplitude matches what you expect from the methods of section 3.9. With these confidence-builders, you are perhaps in a position to trust the numerical solution in the regime where it's indispensable, namely early in the motion where high velocities occur.

Now, check to see if this model produces a trajectory θ(t) which looks like your experimental data. Clearly, you'll have to try various values of the constant, *b'*, to see if you can get a 'best match'. It may happen that the assumption of a single *b'*-value, i.e., a pure v^2-law of damping, is too optimistic. But whatever comes out, you will at least have had a valuable experience in numerical modeling with differential equations, in a case where non-linearity is important, and where you can get good experimental data.

3.10.5 Damping by self-induction

You've seen that a moving conductor (like your copper disc) in the vicinity of a stationary magnet (like your magnetic dampers) represents an energy-loss mechanism, which can be blamed on 'eddy currents'. What about the opposite case, of a moving magnet in the vicinity of a stationary conductor? Here's a way to test that kind of interaction as well.

For this investigation, hand-pump the Oscillator, and use the angular-position transducer's output signal as a monitor of its oscillations. The resulting oscillations can be analyzed to give not only the 'damped frequency', but also the decay constant, of the oscillations. For this investigation, remove entirely the magnetic dampers, to reduce damping to bare minimum -- and recall that small-amplitude oscillations seem to give the smallest possible losses. Now you're ready to search, with considerable sensitivity, for any new loss mechanism.

To do that, let the freely oscillating rotor have its magnets induce an emf in the Helmholtz coil, and confirm that emf is there by looking at the usual velocity outputs of the oscillator. Now flip the toggle switch on the front panel, to bring this emf out of the coils directly to the grey banana-plug outputs -- no filtering action needed in this investigation. Measure the damping constant of the oscillator with those grey terminals left *open*. Next, measure the damping constant of the oscillator with those grey terminals *shorted* instead. You can go on to see how the damping constant of the oscillator depends on the value of R_{ext}, the value of some general resistor you attach across these terminals -- you've already seen the limiting cases of $R_{ext} = \infty$ and $R_{ext} = 0$.

Work out the theory of this effect. To do so, make a model for the emf generated by the magnets-in-coil system -- you've worked this out in section 1.4. The 'calibration constant' of this model is not an arbitrary number -- it's related to other, more easily measurable parameters in section 2.6. Now, suppose that this emf is generated in a coil of some non-zero resistance R_{int} (which is 7-8 Ω), and that it drives a current in a circuit completed by another resistor, R_{ext}. (You can safely ignore the inductance of this circuit, since $\omega L << R_{int} + R_{ext}$ at the relevant frequency ω.)

This gives a model for the coil current, and it flows with consequences. In particular, this current generates a torque, according to the same model you worked out in section 2.3. So put that torque into the differential equation for the motion of the rotor, and show that you still get a homogeneous differential equation, but now with a new loss mechanism added to any intrinsic losses. Work out the damping constant predicted by this differential equation, and notice that there are no unknown parameters in your model, just the value of R_{ext} you choose to complete the circuit. See if your measured damping data agree with the model.

Incidental question -- can you now understand why the bobbins for the Helmholtz coils have been built out of nylon instead of (say) aluminum?

4 Driven Oscillations

If you've worked though parts 1 - 3 of this manual, you'll have seen lots of oscillatory behavior of your Torsional Oscillator, but all of those oscillations were *transient*. That is, you excited the system, by one means or another, only at the beginning of the motion, and then let the system evolve on its own, without further excitation. Now it's time to consider continually-driven oscillations, in which there's an ongoing, time-dependent, torque applied to the system.

4.0 Applying the simplest drive

These investigations make use of the Helmholtz-coil system as a driver, so you have to give up using it as a velocity sensor. The torques are most easily modeled if you stay in the small-angle regime, to evade some complications you studied in section 2.2. The driven-oscillator investigations are simplest if you use a small amount of v^1-law of magnetic damping. Finally, for reasons not immediately obvious, the simplest drive torques are of very special form: they are *sinusoidal* in time.

You will need some signal generator or other electronic source of sinusoidal voltages, capable of having its frequency conveniently and continuously adjusted in the range 0.1 - 10 Hz. If your generator has a 50-Ω output impedance, it can be directly connected to the Helmholtz coils of the Torsional Oscillator, and it'll drive a current through the coils. An amplitude setting of 5 Volts would drive a current of amplitude 0.1 A though a short circuit, which is about what the coils 'look like' to such a generator. That will be ample to drive oscillations of adequate size. Finally, it's ideal to have a look at the 'drive signal' and the oscillator's 'response' simultaneously. A 2-channel oscilloscope or equivalent is a great way to do so. If you want one channel to display the drive waveform, you can route the generator current first through the coils, and then (via a jumper) through the 1.0-Ω resistor in your Oscillator. Across that resistor you'll get a voltage signal which is a surrogate for the drive current, with a known scale factor.

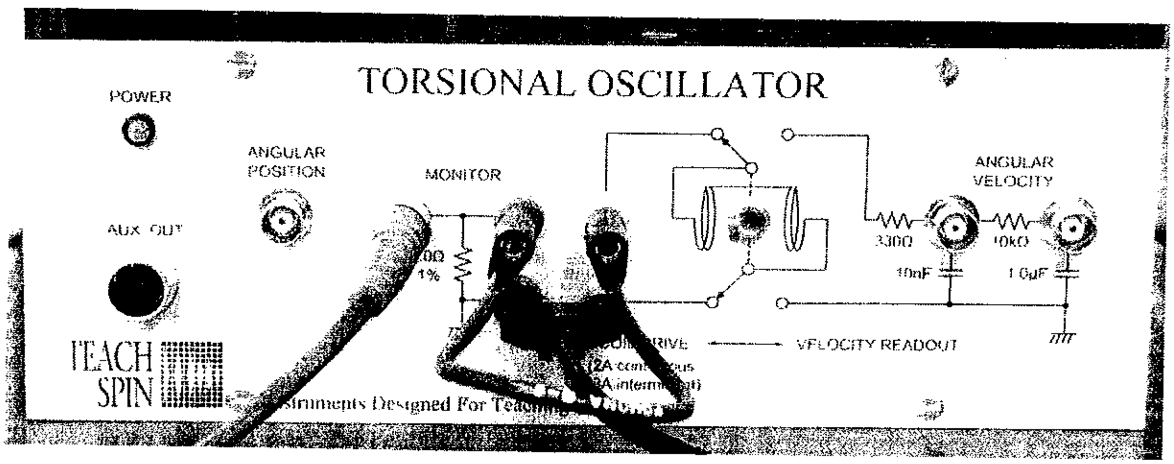

Fig. 4.0: One way to send a generator signal through coils and 1.0 Ω resistor

You will probably want to trigger your 'scope on the drive signal, and now you should see cause and effect, drive and response, quite clearly. Before making any quantitative measurements, note a few qualitative facts:

- after any change in the settings, there's a time delay before you get to a 'steady-state response';
- during that time delay, of 'transient behavior', life is complicated, but the steady-state response is rather simple;
- if the damping is rather larger, then the delay until you reach 'steady state' is rather smaller;
- once you're in steady-state, the response to a sinusoidal drive (of any amplitude, and any frequency) is *also* sinusoidal
- the steady-state sinusoidal response is *not* at the oscillator's 'natural frequency', nor even at its 'damped frequency', but at the *generator's* frequency (and how can you tell this, visually?).

All of these are lessons of enormous generality, and apply across the board, in all sorts of areas and examples of physics, and require only the assumption, or approximation, of a 'linear system'. So learn these qualitative lessons well, as they can give you intuitions of wide utility.

4.1 Resonance in driven oscillations

This section takes up the most 'glamorous' feature of driven oscillations in a damped simple harmonic oscillator -- the phenomenon of *resonance*. To perform these investigations, use the oscillator driven as described in section 4.0, and with a moderate level of magnetic damping in place. You might want to aim for a 'Q' of about 8, and you may want to check section 3.3 for a quick way to estimate the Q of an oscillator. Finally, pick a rather small level of the drive current, setting its amplitude to a value which (applied at zero frequency) would give the rotor a deflection of perhaps 0.05 radians.

Now you're ready to see resonance -- you'll want to monitor the 'drive' and 'response' waveforms in real time. The goal is to make the frequency of the drive waveform the independent variable, holding everything else (particularly, the *amplitude* of the drive) constant. The dependent variable is the amplitude of the steady-state response. Notice that after any change in the drive conditions, you'll have to wait for several Q's of cycles of the drive for the response to settle to steady state -- this wait is a motivation for not making the Q of the system much higher.

Make a plot that shows the amplitude of steady-state response as a function of drive frequency, taking data points at a density suitable to the physics you're investigating. You might take points of higher density in the 'resonant region', and it's your job to characterize not only that region, but also the below-resonance and above-resonance regions.

You should see a plateau, a peak, and a drop-off in various regions of frequency space, and you'll learn to characterize all three. Locate in frequency the peak of the resonance, *and* those points where the response is smaller, by a factor of

$$1/\sqrt{2} = 0.707$$

than the response at its peak. (That's because the *square* of the oscillator's response is down to *half* its maximum value at these points.) You might label the resonance by three frequencies, $f_<$, f_m, and $f_>$, for the lower 'half-maximum', the peak, and the higher 'half-maximum' points. Form the quotient

$$f_m/(f_> - f_<)$$

which is a dimensionless measure of the relative narrowness of the resonance, and which is expected to be quite nearly equal to the Q of the oscillator.. Before changing anything else about the oscillator, disconnect the drive, pump it up to moderate amplitude by hand, and acquire a record of its freely-decaying damped oscillations, since this data gives another, and instrumentally rather distinct, method for finding the Q -- see section 3.3 for details.

Now change the damping to a larger value, perhaps aiming for a Q more like 4 than 8. Repeat the scan over the drive frequency, and plot your data on the same graph you created before. Label the data you obtained under more-damped, vs. less-damped, conditions. You are ready to cure yourself of a very common illusion -- notice that decreasing the damping does *not* really 'narrow the peak' in the usual sense. Said another way, lowering the damping everywhere *raises*, and nowhere *lowers*, the response. Nevertheless, compute the quotient

$$f_m/(f_> - f_<)$$

for the data taken at greater damping, and find that this measure of the peak's narrowness has in fact decreased.

What you've done is to survey a resonant peak, and to locate in the traditional ways the peak's center and also its 'width' -- here, adopting the 'full width at half-maximum power' criterion. You've also seen that the peak response of the oscillator occurs near a characteristic frequency, which you should recognize as familiar. Since you know how to use brass quadrants to change the oscillator's properties, you should be ready to confirm that the resonant peak's location can be moved too. You've also seen a second manifestation of the 'Q' of an oscillator, expressed this time in controlling the width of the resonance.

For yet another manifestation of the Q, look at the ratio

(response's amplitude at peak) / (response's amplitude near zero frequency)

which might be called the 'amplification factor' obtained at resonance. It too should be approximately given by the Q-factor of the oscillator.

4.2 Phase shift in driven oscillations

The data on resonance that you've taken in section 4.1 can all be modeled theoretically, but before going on to that exercise (in section 4.3), you ought here to learn to recognize a less-well-known, but very useful feature that also occurs in the neighborhood of resonance. For taking this data, you might want to use an oscillator with Q set to 4 or even lower -- that's to keep the glamorous 'amplification' phenomenon at resonance from distracting you from this new effect.

The effect you want to look at is the *phase shift* between the drive waveform and the steady-state response waveform. To see the meaning of this, it's best to pick a drive frequency well below resonance, and to get a dual-trace 'scope display of the drive, and the response. You're looking 'in the DC limit' in which you can think of the drive as taking on a succession of independent static values of drive current, and the oscillator responding with a succession of resulting static values of response. In other words, the response waveform ought to be (very nearly) in phase with the drive. (If it appears 180° out of phase, use the 'invert' function on your 'scope to view the in-phase behavior described above.)

You might record the drive and the response in some way that lets you see that there is a systematic, if small, time delay between the drive and the response. You can measure this time delay between a zero-crossing of the drive and the next zero-crossing of the response, and then compute the official value of the phase shift from the quotient

$$(\text{phase shift}) / (2\pi \text{ radians}) = (\text{time delay}) / (\text{period of the drive waveform})$$

It's also conventional to convert the phase shift to degrees. Of course there are lots of other ways to compute this phase shift, for example from least-squares fits or other transformations of the data waveforms.

Once you've learned to measure phase shift by your favorite method, take the data needed to find the dependence of the phase shift on the drive frequency, and make a plot that shows the phase shift's variation with frequency. The phase shift ought to do something relatively dramatic near a particular frequency, whose value you ought to recognize. That dramatic behavior ought to become *more* dramatic if you lower the damping, i.e., raise the oscillator's Q. For at least one setting of the Q, find the locations in frequency of the 45°, the 90°, and the 135° phase-shift points. That's because the theory claims that the location of the 90° phase-shift point *is* the natural frequency of the oscillator, and that the quotient

$$f_{90}/(f_{135} - f_{45})$$

will give yet another measure of the Q of the oscillator.

Precisely because the phase shift is going rapidly through the value of 90° at resonance (whereas the amplification factor is a maximum there, and therefore locally independent of frequency to first order), the location of the 90° phase-shift point is a very useful operational *definition* of the location of the oscillator's resonant frequency.

4.3 The transfer function of an oscillator

You've now measured the amplitude, and the phase, responses of your Torsional Oscillator, as driven by a sinusoidal torque. These turn out to be two facets of an amazingly general property of any driven 'linear time-invariant' system, called the transfer function. The claim is that such a system, driven by a stimulus or cause

$$d(t) = D \cos[\omega t]$$

will exhibit a steady-state response, r(t), given by

$$r(t) = D\,|T(\omega)| \cos[\omega t - \phi(\omega)]$$

Notice there are a multitude of claims implicit in these forms: in particular, the response is claimed, under this sinusoidal drive, *also* to be sinusoidal, and of the *same* frequency as the drive, and to have a size *linearly proportional* to the size of the drive. Notice that the constant of proportionality is itself a function of frequency, and it's called 'the magnitude of the transfer function'. The entire response of the system, magnitude and phase, can be encapsulated in a single complex-valued function, called the transfer function, given by

$$T(\omega) = |T(\omega)| \exp[i\,\phi(\omega)]$$

This form of the transfer function presupposes taking the real part of complex functions to describe drive and response, and also assumes the physicist's conventional use of time dependence of the form $\exp(-i\omega t)$. At one level, there is no more to be said -- in the previous two sections, you've measured the only two ingredients of this transfer function in your 'amplitude response' and 'phase response' investigations. You could, of course, lay out graphically the 'locus' of points that $T(\omega)$ traces out in the complex plane, as ω varies from near DC to frequencies well above resonance.

At another level, there is a lot more to be done. Specializing to the particular case of the Torsional Oscillator, there are ways to factor units and dimensions out of the transfer function, leaving behind a dimensionless 'core' form that generalizes to all sorts of other systems in physics. Here's one way to do this:

- you've been thinking of torque as the 'input' to the system, but you have things calibrated well enough that instead of the torque, you could think of the current producing the torque, or even the monitor voltage generated by running this current through the 1.0 Ω resistor. Thus $V_{mon}(t)$ is a voltage surrogate for the drive.

- you've also been thinking of the angular-position coordinate as the 'output' of the system, but again you have a transducer calibrated so that instead of the angle, you could think of the output voltage reflecting it, $V_{pos}(t)$, as the response of the system.

- if your Oscillator is 'linear and time-invariant', it follows that $V_{mon}(t)$ as a drive, and $V_{pos}(t)$ as a response, have a transfer function connecting them -- and it'll be dimensionless, as it maps a voltage to a voltage.

- furthermore, this voltage-to-voltage mapping has a well-defined value in the 'DC limit', i.e., as the frequency approaches zero. In practice, this means at a frequency well below the resonant frequency. There the magnitude of the transfer function can be measured by putting an actual DC value of current into the drive coil, reading a steady V_{mon}, and then seeing what steady value of V_{pos} results. (You might have to deal with the 'DC offsets' implicit in a real apparatus.)

- factoring out as a constant this DC-limiting value, you're left with the core of a transfer function which is not only dimensionless, but which has, in the DC limit, a magnitude of one, and zero phase shift. The behavior of this core transfer function away from zero frequency contains all the interesting frequency dependence.

Now you should be ready to work out the *theory* of your oscillator, and to make a prediction for the transfer function. You'll need to introduce various physical parameters, and set up the appropriate differential equation, and learn why it's an inhomogeneous differential equation, and how to form the 'particular solution' that is the whole of the steady-state response. And then you should work out the dimensionless form of the transfer function, and the form of it with DC-limiting value of 1, and show that the core of it depends on just three things: your only independent variable, the frequency, ω, (or $f = \omega/2\pi$ instead), and just two parameters, constants that characterize your oscillator, namely its 'natural frequency', ω_0, (or $f_0 = \omega_0/2\pi$ instead), and its dimensionless damping coefficient, γ.

Of course, once you have a theoretical prediction of this form, you can find the ω_0 and γ values that best describe your data, by fitting to your amplitude- and phase-response data. Or, you can be more ambitious, and extract ω_0 and γ from the damped-oscillatory-decay data of the sort you modeled in section 3.2. Such data will *also* produce measured values for ω_0 and γ, and these would allow you to make a (zero-free-parameters!) *prediction* of both the magnitude and the phase of the transfer function, before you ever turned on your sinusoidal generator. See if you can make a single pair of numbers, ω_0 and γ, fully describe both the undriven, free-decay data, *and* the magnitude and the phase pieces of the transfer function of the sinusoidally driven oscillator.

4.4 Non-sinusoidal periodic drive

You might wonder about the claim that the transfer function fully describes the response of a system, when after all it seems merely to predict the response to a periodic and sinusoidal drive. What about a drive function that's periodic, but not sinusoidal? (Think of changing your signal generator from sine-wave to triangle-wave excitation.) Or what about a drive function that's not even periodic in time? (Think of some random time series, like a temperature-vs.-time graph, as a signal -- putting it into your Oscillator as a stimulus would surely give *some* response.)

The true power of the claim of linearity shows up in these cases via the wonders of Fourier's theorem. Here's how it works -- just as all molecules are made from a rather short list of atoms, so any and all drive signals can be written as the sum of a weighted list of sinusoids, i.e., they have Fourier-series, or Fourier-integral, representations. Now the big payoff of linearity is the Principle of Superposition:

- the response to a drive which can be written as a sum of functions is given by the *sum* of the responses to the individual terms in the drive function.

And the individual terms in the drive, in this view, are sinusoids, each of a particular frequency. The transfer function, evaluated for magnitude and phase information at that frequency, gives the predicted response to that term in the Fourier sum. Similarly for every term; and the resulting response is just the sum of all those individual responses.

This takes some practice fully to appreciate, but you should start on the empirical side. You might by this time think it's obvious and natural that the response to a sinusoidal drive is a sinusoidal response. But check for yourself that the responses to a triangle-wave drive is *not* a triangle-wave response! In particular, pick a moderate-Q system (say, a Q of 5 to 10) and excite it with a triangle-wave drive, with a frequency picked to lie close to the system's natural frequency. What do you see? Or, for more drama, excite the same system with a square wave of the same frequency -- what do you get? More curiously still, excite the system with a square wave of one-*third* the previous choice of frequency -- what do you get?

The answer to all of these questions, certainly in qualitative terms, is best found from a Fourier-series point of view. The method works in quantitative detail, too, if you care to work that out. But it's indispensable to have in your mental tool-kit the 'frequency-domain view', to complement the 'time-domain view' you've found natural so far. What *is* a triangle wave, in the frequency-domain view? What are its terms, and what can you say qualitatively about the response to those terms, taken individually? Which will dominate? What shape will that response have? You have before you the tools that will make it natural to think in this new way.

4.5 Time- and frequency-domain views

If you have played with non-sinusoidal drive of your Torsional Oscillator, and have begun to learn all the insight that can be gained using a frequency-domain point of view, you might be willing to give up the time-domain point of view you first found natural. Here's an exercise to show that the two viewpoints are actually complementary, and co-exist fruitfully, in making the behavior of a system understandable.

Set up your Torsional Oscillator as a moderate-Q system, perhaps using a Q of order 10. Again, excite it with a non-sinusoidal but periodic drive function -- you might try a 'square wave', chosen because of its considerable harmonic content (to use a frequency-domain vocabulary). This time, look not at the detailed waveform-in-time of the response, but simply measure the rms value of the response function. (In practice, this means exciting the system, waiting for steady-state conditions to obtain, recording the response function over one or more full cycles of the drive, and then forming the root-mean-square value of that response.) Now the question is -- keeping the drive's size constant, and varying its frequency, how does that rms value of the response vary with that choice of drive frequency?

In taking the data, you will want to pick excitation frequencies thoughtfully. Your intuition should suggest you need to pay attention to the three regimes of below-, near-, and above-resonance. But in this case, you'll find you need to take extra data where the drive frequency lies near *sub*-multiples of the system's resonant frequency. A frequency-domain view will help you understand why -- and why the 1/3, 1/5, . . sub-multiples repay more effort than the 1/2, 1/4, . . . sub-multiples.

After you get the idea of what will turn out to be a complicated graph, you should be able to 'narrate' the result using frequency-domain ideas. But just before you convert entirely to this viewpoint, you might look past the rms value to the response waveform itself, viewing it in the time domain again. Since your complicated graph taught you there was a lot of structure in the low-frequency regime, set the drive to a rather low frequency, and see what occurs that's so special near those odd sub-multiples of the resonant frequency. Go lower still in frequency, and suddenly you'll see things from another viewpoint, one best described as a series of excitations of the *step response* of your Oscillator. That puts you back in a time-domain vocabulary, and teaches you that the value of having *both* time- and frequency-domain viewpoints as part of your mental equipment.

4.6 Transient behaviors

You've been looking at the behavior of a driven oscillator for some time now, and you have fixated (or been led to fixate) on the steady-state response. There are good reasons for this -- in particular, that response is independent of those pesky initial conditions of the system. But there's lots of good physics in the 'transient behavior' of the system, all the response that happens before the system settles to its steady-state.

In studying such behavior mathematically, you have perhaps been told to solve the problem in two parts. First, you find the general solution to the un-driven problem: the homogeneous differential equation. Then, you find the particular solution to the driven problem -- the *in*homogeneous differential equation. Then you write the whole solution as the sum of these two parts, and finally you apply initial conditions to that whole sum, to fix the last of the unknown constants. Whew! But the implication is that the solution, for a good while after the launch from initial conditions, is the sum of two parts, which oscillate at possibly different frequencies. The particular solution oscillates only at the drive frequency, and it has constant amplitude. But the general solution, which is important until it decays away, will oscillate at the 'damped frequency'

$$\omega_d = \omega_0 \sqrt{(1 - \gamma^2)}$$

So the total solution will be the sum of two terms of distinct frequency -- and that sum should therefore display 'beat phenomena'.

To move from the mathematical to the physical world is now easy for you. Adjust your Torsional Oscillator for rather low damping, so that the Q might be of order 20 to 40. Stop the rotor movement by disconnecting the drive signal and damping it carefully by hand. You'll want a way to connect the drive suddenly at your t=0 point. Choose a drive frequency which lies in the vicinity of the system's 'natural frequency', and choose quite a small amplitude -- you'll be driving a high-Q system near resonance, and the steady-state response is subject to a large 'amplification factor' because of resonance.

Start by setting the drive frequency right at your best estimate of the system's damped frequency, and record the whole history of the oscillator's response. You should see an oscillation of growing amplitude -- growing linearly at first, and later leveling off at a steady-state amplitude. Once you've understood that, both physically and mathematically, try the same experiment, only this time choosing a drive frequency that is (say) 5% *off* from the system's natural frequency. You should see short-term behavior that is nearly identical to the previous case, and then moderate-term behavior that shows dramatic 'beats', before you finally get to the steady-state behavior. What sets the period of these 'beats'? What does it take for them to have maximal amplitude? Repeat with a drive frequency that's 2%, or 1%, or -2%, etc. off from the natural frequency. In each case, look at the short-, moderate-, and long-term behavior of the system to see if your mental vocabulary contains the concepts needed to explain what's changing, and what's not. And don't forget to tear your attention away, from time to time, from your data-recording system, to have a look at your

Oscillator's copper disc, to see it undergo the complicated motion that a driven system can (transiently) possess.

4.7 Projects

Here is a collection of projects generally unified under the theme of driven oscillators.

4.7.1 The oscillator as a low-pass filter

If you've completed section 4.3, you know how to compute (in magnitude and phase) the theoretical transfer function for your Torsional Oscillator. While the high-Q version of the transfer function displays glamorous features around the resonant frequency, there is merit in plotting the transfer function, in magnitude and phase, in some rather low-Q situations. In fact, certain choices of Q, or the damping

$$\gamma = 1/(2Q) ,$$

give responses famous enough to have been given names.

Work out an analytic expression for the *magnitude* of the transfer function, and plot it for a variety of rather low Q values (i.e., for some rather *large* γ values). You should notice that the choice

$$\gamma = 1/\sqrt{2}$$

has a property that renders it special. In particular, you get a transfer function whose magnitude is 'maximally flat' in the low-frequency regime. Of course it shares features with the transfer function for any γ-value: it has a value going to 1 in the low-frequency limit, and it has a drop-off of the form $(f_0/f)^2$ in the high-frequency limit. The $\gamma = 1/\sqrt{2}$ value gives the special feature of low frequency response of 'maximal flatness', and this choice is named 'Butterworth response'. Your system's behavior can be viewed as a low-pass filter, in that low frequencies at the drive input are 'passed' to the output, while high frequencies are (somewhat) blocked from appearing at the output.

Next, work out an analytic expression for the *phase shift* of the transfer function, and plot it too for a variety of rather low Q values (ie. for some rather large γ values). You should notice that the choice

$$\gamma = (\sqrt{3})/2$$

has a property that renders it special. In particular, you get a transfer function whose phase shift is 'maximally linear' in the low-frequency regime. This turns out to render its 'group delay' for a pulsed waveform as frequency-independent as possible, and this choice is called 'Bessel response'.

Note that neither of these responses is critically damped, since both require the choice $\gamma < 1$. Since both Butterworth and Bessel response are *under*-damped, it follows that their *step* responses will show some overshoot. You can work this out analytically, and display it graphically too. But as you learn how to recognize the effects of varying the damping, you have the opportunity to set up your actual Oscillator to a chosen γ-value, and to get it to the point of displaying actual Butterworth, or a

Bessel, filter response. Once you have it set up in your chosen way, it's up to you what to do with it -- you can send in sinusoids, and confirm the filtering action via the amplitude response; or you can send in step-functions, and look for the overshoot in the time-domain response; or you can find a way to make 'wave packets' and measure their phase, or their group, delay in passing through.

This is just an introduction to the well-developed theory of filter design, and you've now worked through 'two-pole filters', which is the most you can do with a single oscillator as a filter element. Can you think of a way to turn your Oscillator into a narrow-band *band-pass* filter?

4.7.2 Life beyond linearity

Everything thus far in section 4 has had you using the magnetic dampers, whose v^1-law of damping offers the remarkable mathematical property of linearity. One of the many consequences of linearity is that scaling the input will scale the output in the same proportions. Within limits, your Oscillator, under magnetic damping, shares this property. In this section, you'll get to see some evidence of *departures* from this property of linearity, under magnetic and other damping laws.

You can see the consequences of linearity in some of the simplest properties of the oscillator. The simplest of all is to drive the system with a sinusoid, of frequency chosen to lay on-resonance, and with amplitude to be adjusted. Then you need only look at the output, which (under the assumption of linearity) will also be a sinusoid, with some output amplitude. A plot of output amplitude as a function of input amplitude is a test of linearity as a system property -- in particular, a 'linear system' will give a straight line for this plot.

Try setting up your Oscillator with moderate Q, of perhaps 5 or so, and use one or another criterion to find the resonant frequency and operate there. Try a wide variety of choices for the amplitude of the drive, and for each choice, wait for steady-state operation, and record the amplitudes of input and output, i.e., drive and response, waveforms. Make a plot, and find a <u>departure</u> from a straight line. Why? Look back to section 2.3 and realize that your torque drive gives torque that is proportional to the drive current, but with a proportionality 'constant' that is actually constant only for small-amplitude motion of your oscillator. Nevertheless, you should find a small-enough-drive regime in which you *do* see a straight-line plot, as system linearity predicts.

Now change to either the v^0- or v^2-law of damping, and life will get more complicated. For these cases of damping, it's not even obvious how to *define* the 'resonant frequency', since the location of maximum response can actually vary with choice of drive amplitude. But you can make *some* sensible choice that will give some version of resonance in some regime of drive, and then at that fixed frequency, you can again plot amplitude of response vs. amplitude of drive. You will see very different plots than displayed by your former data! In particular, for the v^0-law of sliding friction, there ought to be a threshold level of drive, below which you get no response at all -- why? And what happens above this threshold? For the v^2-law approximated by fluid friction, you will get another quite distinct plot, with approximate power-law dependence. Can you predict why, and what exponent you expect in the power law? You might use an 'energy budget' that applies to steady-state operation: averaged over a full cycle, the work done by the drive and the work done by friction, have to be *equal* in steady-state operation.

4.7.3 Intermodulation in non-linear systems

You've become used to the idea that if you drive a system sinusoidally at frequency, f, out will come a response which will *also* be sinusoidal, and of frequency, f, too. This is not automatically a property of all driven systems! Or, consider another property of linear systems: if the input is a superposition of two sinusoids, at frequencies, f_1 and f_2, out will come another superposition of sinusoids, containing (only) the same two frequencies, f_1 and f_2. In this case, of course, the two sinusoidal components might well have experienced different magnitude response, and different phase shifts, but still there are only two frequencies represented in the output.

One of the most characteristic features of *non*-linear response is 'intermodulation', and the 'two-tone' input waveform containing frequencies, f_1 and f_2, is just the way to diagnose such systems. The test is to get a frequency-domain view of the output waveform, and in particular to look for any frequencies other than f_1 and f_2 to appear. Under conditions of even rather weak non-linearity, you should expect to see 'intermodulation distortion', with the appearance of frequencies including $f_1 + f_2$, $| f_1 - f_2 |$, and other small-whole-number combinations too.

You might test this first on your Oscillator working as a rather low-Q system under magnetic damping. The resonant peak will be broad, so you can pick two frequencies that both lie not too far from the peak. Give them enough separation such that the difference, $| f_1 - f_2 |$, will not be too small, and choose a low level of drive (to evade the known non-linearity implicit in large-angle operation of your torque drive). Now you'll need to wait for steady-state operation (i.e., for all the transients related to initial conditions to die away), and you'll need to get a spectral view of the output via the Fourier transform. You're looking for just two peaks to appear, and in particular, you can check to what degree (i.e., by how many decibels, on the traditional logarithmic scale) the expected 'intermodulation tones' are suppressed.

Now change to a higher level of drive, where you do expect some degree of non-linearity, and see if you can detect it by this test. Once you've confirmed by these two 'control group' experiments that you can get proper negative, and positive, results of the test, it's time to change to a different law of damping, and see if intermodulation distortion is present, as theory predicts it should be. As you raise the level of drive, you might see *lots* of frequencies, of the form

$$|m f_1 \pm n f_2|$$

for integers, m and n, appearing in your Fourier spectrum. Along the way, you'll learn that for high spectral resolution in your frequency-domain view, you need to take records of output waveform that are of long duration in the time domain.

4.7.4 The Kramers-Kronig relationships

You've learned about transfer functions: that the complex-valued transfer function conveys information on both the amplitude response, and the phase-shift response, of the system it describes. Now you might imagine that you can build a system with a tailored-to-choice magnitude-response function, and also with a tailored-to-choice phase-response function. For example, you might want to build a filter that:

- has amplitude response of one in a chosen band of frequencies (but zero response elsewhere), and further
- has no pesky phase shifts at all.

But such specifications are not realizable for generic systems! In particular, there are very powerful statements that can be made about any system that is linear and *causal* -- that is, for which effect can only follow (not precede) the cause. For a causal system, in fact, the transfer function is so tightly constrained that knowing the magnitude response fully fixes the form of the phase-shift response, and conversely too.

The name for these tight connections is the 'Kramers-Kronig relationships', and in advanced texts you can find, in closed form, the expressions that give the phase response (if the amplitude response is known), or conversely give the amplitude response (if the phase response is known). You will learn a great deal about the mathematics of complex variables if you follow the standard derivations and a good deal about your oscillator system if you start to ask just how you would use the Kramers-Kronig relationships operationally.

For starters, you might work with a paper system, i.e., a torsional oscillator entirely described by a simple model of a second-order differential equation. You'll get an exact and closed-form result for the complex-valued transfer function, and you can temporarily pretend you know only the magnitude of that function. Get a bit more realistic, and suppose you knew the amplitude-response function at only some finite list of frequency values. Can you use this finite set of information to execute the actual integral that gives, according to the K-K relations, the phase response? In this paper example, you can of course check those results against the predicted phase response of the system.

Life gets harder still with actual experimental data. Here, you not only are limited to some finite list of values in a finite range of frequencies, but you are also limited to some finite level of precision in your ability to measure the amplitude response at any of them. So it's actually a rather difficult problem to execute the K-K program in practice, and you can look into the research literature for suggestions. But practicality aside, there's so much *romance* in the far-reaching consequences of a stipulation so apparently obvious as causality.

5 Advanced Topics

This section of the manual contains some advanced-project topics that are not so naturally associated with the sections you've seen so far. Some of them also involve the use of some additional electronic equipment. Any of them could take you deep into a research topic, and all of them are covered in these notes in rather less detail than you've seen heretofore.

5.0 Overview of topics

Here are the 'abstracts' for the advanced topics -- more complete descriptions follow:

1. Drive by noise waveforms

   You're used to treating electronic noise as a nuisance, or an outright enemy -- but here's a chance to see its usefulness, and learn the basis of 'Fourier transform spectroscopy' at the same time.

2. Feedback and its effects

   What happens when you take the angular-position signal, process it according to some recipe, and 'feed it back' into the drive coil of your Oscillator? You can change the properties of the oscillator by that feedback, varying either the natural frequency or the damping.

3. Building a 'torsional clock'

   If you've worked on feedback in section 5.2, and seen that you could change the damping of the Oscillator, you might have wondered if you can change the damping all the way down to *zero* -- and what happens then?

4. Aligning the magnet-in-coil differently

   Everything you've done thus far had the magnets on the rotor, at its equilibrium position, come out perpendicular to the coils' axis. How can you change this to the parallel configuration, and what's the usefulness of this new geometry?

5. The quartic oscillator

   The elastic behavior of the torsion fiber has given you a torque linear in displacement, which corresponds to an elastic potential-energy function that's quadratic in displacement. How can you change this to a *quartic* potential, and what is the result?

6. Parametric drive of an oscillator

You're used to driving an oscillator at its own 'natural frequency' to get resonant behavior. Under what circumstances can a drive at *half* the natural frequency nevertheless pump up an oscillator?

7. Coupled oscillators

If you have the luxury of two Torsional Oscillators, you can couple them magnetically, and see the surprising phenomena that result. There are plenty of qualitative observations, and quantitative measurements, that will enable you to learn the very generally applicable lore of coupled oscillators.

5.1 Drive by noise waveforms

In sections 4.1 and 4.2, you learned how to measure the amplitude, and the phase, response of your driven oscillating system. In both cases, the independent variable of your plots was your choice of the frequency of sinusoidal drive you put into your system. In both cases, you took the data by putting in one frequency at a time, and perhaps you got bored by the time you had measured the tenth or twentieth point. Here's a way to get *all those points at once*, and lots more points besides. The method is applicable only to linear systems, but it's wonderfully efficient where it works.

The secret is to take advantage of linearity in a serious way. Since the system's response to a sum-of-inputs is just the sum of the responses to the inputs applied individually, the idea is to subject the input of the system to a waveform that contains *all* frequencies at once. You can even ask for all frequencies to be 'equally represented', and the name for one version of that waveform is 'white noise' ('white' by analogy to white light, containing as it does all the frequencies of visible light). In particular, for this investigation you need some sort of white-noise generator, in place of your usual signal generator limited to sine-square-triangle waves.

Of course you don't really want <u>all</u> frequencies, since you don't need to send gamma rays into your system. In fact, for these experiments on your Torsional Oscillator, all you need is decent coverage of the 0-10 Hz range, since all the interesting response of your oscillator appears in that range. There are lots of ways to get noise that is effectively white over such a range, and you can pick any of them.

Pick a strength of the input signal which is reasonable -- one method is to measure the rms value of the white noise, and make it comparable to that of the sinusoidal drives you've used before. Set the Q of your system to order 5, such that the magnitude response of the transfer function ought to have a well-defined, but not too narrow, peak. Finally, in this experiment you'll need (possibly simultaneously-acquired) time records of both input and output waveforms, and you'll need to get frequency-domain views of both waveforms.

The spectrum of your <u>input</u> waveform ought to confirm that (over some frequency range) your signal is indeed 'white'. Or, if it's not, you will have a measure of how much power it contains at each of the frequency points at which you diagnose it. The time-domain record of your system's <u>output</u> will look like *garbage* -- in fact, you're putting noise in, so the best you can hope for is noise out. But this will be amazingly informative 'garbage', if you get a frequency-domain view of the output. In particular, if the input spectrum is white, the output's spectrum will display, all in one go, the transfer function of the oscillator! If you can get the complex values of the Fourier coefficients of both the input and the output waveforms, acquired simultaneously, you can go on to get the phase shift of each frequency component, and similarly get the phase-shift plot not point-by-point, but for all frequencies in one investigation.

This very powerful effect, of exciting a system with all frequencies at once, and then separating the frequencies after the fact, is powerfully put to use in Fourier-transform spectroscopy and lots of other places. You can learn a great deal of transferable knowledge about these Fourier-based techniques even on so simple a system as your oscillator.

5.2 Feedback and its effects

In this section you'll see some of the modifications you can make in the behavior of the Oscillator by external electronic changes, instead of internal mechanical changes. The general technique is feedback, in which a signal is taken out of the mechanical system, electronically modified, and fed back into the system. To perform these experiments, you'll need some external 'bread-boarding' analog-electronics capabilities.

Start by getting the angular-position signal, $V_{pos}(t)$, out to an electronic arena in which you can modify, or amplify, it. In this same arena, you want electronics that can drive the Helmholtz coils, not with some signal generator's output, but with a signal derived from $V_{pos}(t)$ itself. Possibly the best way to drive those coils is to build a voltage-to-current converter, which has its output current pass through the Helmholtz coils. Model your circuit according to

$$i_{coil} = V_{V\text{-to-}i} / R$$

where $V_{V\text{-to-}i}$ is the input voltage to your coil-driving circuit, and R is some quantity (of dimensions resistance) characterizing your circuit.

Now write the differential equation describing the motion of the oscillator, under conditions of some generic function, $V_{V\text{-to-}i}(t)$. But go on to assume, not a generic value of $V_{V\text{-to-}i}$, but instead a value *proportional* to $V_{pos}(t)$, the angular-position output of your oscillator. You get to choose, by actual electronic implementation on your breadboard, the *sign* and the value of this proportionality constant. And since $V_{pos}(t)$ is itself proportional to $\theta(t)$ with a known constant of proportionality, you should now be able to get a homogeneous differential equation, i.e., one involving $\theta(t)$ and its derivatives, but no unknown or external function. Your system is now electro-mechanical in character, but it's autonomous, with no connection to the outside world, acting under 'proportional feedback'.

Show that your differential equation corresponds to that of a modified simple harmonic oscillator. In particular, the 'natural frequency', ω_0, will have been changed, by an amount which you can calculate with no unknown parameters. Because of the change in ω_0, there will also be changes in the damping parameter, γ, also by a calculable amount. By your choices of circuit connections, you can make this new natural frequency larger than, or smaller than, the former value. Which direction of change has the effect of making the oscillator more, or less, damped? Can you get data confirming both effects?

Your voltage-to-current converter, driving the coils, will have some upper limit to the size of currents it can deliver, and this will limit the size of oscillations your system can undergo and still fit your model. But within these small-oscillation limits, you can still measure the frequency and damping of small oscillations, and gain confidence in your model.

The modification of the behavior of systems in general, and instruments in particular, by external electronic circuits and feedback is a very useful general technique. Suppose you wanted to achieve the effect of a much smaller torsion constant κ for your fiber, so as to give your torsion apparatus some extra sensitivity to small torques -- can you achieve that with feedback?

5.3 Building a 'torsional clock'

In the previous section, you've seen some of the uses of feedback in modifying the mechanical behavior of your Oscillator by the use of external electronics, and feedback. Here's another exercise along that line, with even more glamorous results.

Once again, you will use the position signal, $V_{pos}(t)$, coming out of your Oscillator as the input to some external electronic breadboard, and once again you will use a voltage-to-current converter, or the equivalent, as a way to send currents into the drive coils of your Oscillator. What's different, compared to the previous section, is the interposition of a extra stage of processing between $V_{pos}(t)$ and the voltage, $V_{V\text{-to-}i}$, that you send into your voltage-to-current converter. The stage of processing that will generate the dramatic new results is to put either an electronic differentiator, or an electronic integrator, in this intermediate position.

Note that you can build analog-electronic realizations of both these circuits, with very nearly ideal performance from the integrator, and adequate low-frequency behavior from the differentiator. Note that you can build the circuits so that they'll behave according to well-defined models -- that is to say, such that they introduce no unknown constants. Notice to that you still have control of the *sign* of the whole feedback term, either by circuit changes, or just by interchanging the two leads at the grey-banana coil connections to the drive coil.

Before you build or test these circuits, try working out the differential equation that will describe the behavior of your electro-mechanical system. With the *differentiator*, you should retain a second-order differential equation, and you should be able to see that in this case, the electronic modifications change the decay constant of the system. Note that you can make the system decay faster, and offer the advantage of faster 'settling time', or decay slower. A system decaying slowly enough will never settle -- so, show on paper that you can reduce the net damping from whatever former damping you had, all the way to zero. What ought to happen then?

Similar analysis of the circuit with the *integrator* will give a third-order differential equation, still linear and solvable in principle. Rather than wade through the details, find out under what conditions this equation has an undamped sinusoid as a solution -- that's a lot easier, mathematically, and it'll lead you to a choice of parameters that you can actually build.

The payoff is rather exciting -- you will have an electro-mechanical system which is autonomous, in which the external electronics are arranged precisely to make up for the mechanical damping that was originally present. Such a system ought to oscillate without damping, and in fact will oscillate at its own natural frequency as a 'clock'. If you measure the period of oscillation electronically (say, with a digital counter), you'll find you have remarkable short-term stability, and amazing sensitivity to small changes in the period caused (for example) by the addition of tiny masses to the rotor.

5.4 Aligning the magnet-in-coil differently

In every investigation thus far, the equilibrium position of the magnets on the rotor shaft has been *perpendicular* to the axis of the Helmholtz coils. This makes the torque of the coils on the rotor a simple function (at least for small angles of deflection), and also allows the coils-in-magnet to serve as a velocity transducer. But in this and the next two sections, that arrangement is changed, to a situation in which the magnets' moment is aligned to lie *along* the coils' axis. The usefulness of this arrangement is found in the next two sections; this section tells you how to achieve this re-alignment.

The basic idea is to loosen the rotor shaft from attachment to the torsion fiber using the wire clamps at the top and bottom of the rotor shaft, to turn the whole rotor by 90°, and then re-tighten those clamps. Here are the details:

- slide the rotor set-up tool (typically stored atop the upper support of the Helmholtz coils) forward, until it's under the rotor;
- use the tensioning knob at the top of the instrument's case to slacken the torsion fiber somewhat, until the rotor rests on the rotor set-up tool;
- loosen the wire-clamping screws at the top and bottom of the rotor (these are the socket-head screws whose axes lie in the *horizontal* plane) by a turn or less;
- slightly loosen the rotor-mounting screws in the wire clamps (these are the socket-head screws whose axes are *vertical*), which will allow the clamps to lose their grip on the fiber;
- rotate the whole rotor structure by 90° (either way) until the permanent magnets' flat faces lie perpendicular (and their magnetic-moment vectors lie *parallel*) to the axis of symmetry of the Helmholtz coils;
- re-tighten the wire-clamping screws on the wire clamps, symmetrically;
- re-tighten the rotor-clamping screws on the wire clamps;
- (if desired) remove all four screws that hold the copper rotor disc to the rotor shaft, and (counter)rotate that disc by (-)90°, so that it's restored to its original orientation (this permits the familiar use of the 'radian protractor scale'), and then re-fasten it with the four screws;
- (if desired) remove all four screws (down near the bottom of the rotor shaft) that hold the rotating central plate of the angular-position transducer, and (counter)rotate that plate by (-) 90°, so that it's restored to its original orientation (this permits the familiar use of the angular position transducer), and then re-fasten it with the four screws;
- re-tension the fiber, slide the rotor set-up tool out of the way, and check to see if you've achieved the desired orientation of the magnets' axes -- if not, iterate the above procedure. (Small deficiencies in orientation can be corrected using the fiber angular adjuster at the top

of the torsion fiber.)

You will see in the next section that there are rather direct checks to see if you have indeed achieved the parallel orientation of magnetic moments, and coil axis, that you seek. If you can achieve alignment to 'eyeball accuracy' by the procedure above, the use of those diagnostics will allow you to achieve optimal alignment.

5.5 The quartic oscillator

The previous section has taught you how to achieve a novel orientation of the rotor magnets' axes along the axis of the Helmholtz coils of your Oscillator, and this section will show you how to diagnose that alignment, and to test it systematically, and finally to apply it to some glamorous new kinds of oscillation.

If you have, mounted on the rotor at elastic equilibrium, a magnet system with axis exactly aligned along the Helmholtz coils' axis, then adding a modest current in the coils will not change the location of the equilibrium position. But clearly, if those axes differ, then the equilibrium position due to elastic and to magnetic interactions acting jointly will be a compromise between minimizing elastic energy and minimizing magnetic energy. The consequences will be most obvious with the use of a small current in the coils, chosen to have a direction that would (by itself) create a *de*-stabilizing effect on the rotor.

If you see a change in the equilibrium location of the rotor with modest sizes of this sort of current, then you have a signal diagnostic of imperfect alignment. You can then use the angular adjuster at the top of the torsion fiber to fine-tune the zero position of the rotor for elastic-only interactions, until you achieve the alignment desired.

To get some quantitative indication of the effects of the combined elastic-plus-magnetic interactions, you can choose a sign and magnitude of the current in the coils, and then hand-excite a small oscillation about the equilibrium position of the rotor. If you plot the square of the angular frequency of oscillation, ω^2, of this motion, against the current, i, you send through the Helmholtz coils, then you should get a linear dependence. (Why?) Make a model of the potential energy of this system (elastic plus magnetic), and use it, or a torque equation, not only to understand why such a linear dependence is observed, but also to understand its slope and intercept in terms of other parameters of your electro-mechanical system.

In particular, for a certain value of current, i_b, (with a sign that makes for *de*-stabilizing magnetic interactions) you can extrapolate to find where your small-oscillation frequency, ω, goes to zero. This is the point at which the quadratic and positive *elastic* contribution to the system's potential energy is 'balanced out' by a *magnetic* and negative contribution. In fact, you should draw models, as a function of the rotor's deflection angle, θ, and the current in the coil, i, of the total potential-energy function of the system, and show that

- for currents in the vicinity of i=0, the potential energy has a quadratic minimum;
- for one sign of currents (taken to be positive), that minimum becomes 'deeper' with current;
- for currents of the *other* sign, headed down to i_b, the potential-energy function gets 'softer';
- for a current of exactly (-)i_b, the quadratic term to the total (elastic plus magnetic) potential energy *vanishes*;

- at that current, the next surviving term is quartic (θ^4) in angle, and *restoring* in character: and
- beyond (-)i_b, the potential's minimum *bifurcates*, leading to two stable equilibria.

You will note that it doesn't pay to go far beyond (-)i_b in current, since you can very easily stray into a regime where the whole rotor will suddenly slew through a large angle, and oscillate with large amplitude about an angular position 180° away from your original position, where the magnets are favorably aligned with the magnetic field of the coils. Even short of that point, you will still discover that in the vicinity of (-)i_b, the system as a whole becomes amazingly sensitive to small displacements (and to the 'memory' effect of hysteresis in the fiber).

Nevertheless, you should be able to be confident about the value of the 'balancing current', and should be able to obtain stable oscillations about the potential-energy minimum you get using that current. But they will be oscillations of a 'quartic oscillator', whose small-displacement behavior is dominated by a potential of the form

$$U(\theta) = q\,\theta^4$$

for some positive (and predictable!) constant, q. Motion in such a potential does indeed have a restoring force, but it's *cubic*, rather than linear, in the displacement. It follows that this motion will *not* be sinusoidal, and will not be isochronous either -- in particular, the period of small oscillations will depend, rather dramatically, on the amplitude of the oscillations.

It's a relatively complicated task to predict the period of those oscillations, even in a pure-quartic potential -- and for amplitudes of any finite size, your potential is not exactly quartic. But the potential is entirely predictable, if you assume a pure-quadratic elastic contribution, a magnetic contribution of the expected form, and operation right at the 'balancing current'. In a well-understood potential, both the motion, and the period of the motion, of the system ought to be entirely predictable by one or another numerical means.

If you can operate *just beyond* the balancing current, you'll have the chance to see motion in a double-well potential. In principle, you could measure the frequency of small oscillations in each of the two wells, each of which ought to be (locally) quadratic. In practice, hysteresis effects will complicate your understanding of this situation. But it's still worth thinking about this situation, since a double-well potential with a small 'barrier' between the two wells is just the sort of system in which dramatic effects like tunneling are predicted to occur quantum-mechanically.

5.6 Parametric drive of an oscillator

Here's another curious capability of your Torsional Oscillator, placed here in the manual since it requires the magnets-along-coil-axis geometry set up in section 5.4 and exploited in section 5.5. You've modeled the interactions of the magnets, so oriented, in the previous section, when you created a static magnetic field using DC currents in the drive coils. Now you're ready to try exciting the oscillator by using *AC* currents in the coils.

Everything in your preparation has led you to expect that this effect will be resonant when the frequency of the drive current matches the 'natural frequency' of the Oscillator. In fact, let's review what happens with ordinary drive of the oscillator, as in section 4.0. For any frequency of the drive, you get *some* response of the oscillator, but you get the most response if you're 'at resonance'. For any amplitude of the drive, you get response of the oscillator, and the steady-state response is linearly proportional to the drive. If you start with a quiescent oscillator, and drive 'on resonance', the response grows linearly in time at first, before leveling off toward the steady-state amplitude.

Now here's the contrast. In the 'parametric drive' that you're about to undertake, you will indeed get the most response if you're 'on resonance', but here, 'on resonance' will require a drive frequency *double* that of the 'natural frequency' of the oscillator. If you're off resonance a bit, you'll get less response, but if you're off by *too much*, you'll get no steady-state response at all. On the matter of drive amplitude, you'll find there's a *threshold* amplitude, below which oscillations will only decay; and above threshold, you'll find oscillations do grow, but exponentially in time. Only exactly at threshold will you get what you thought was normal: steady oscillations of stable amplitude.

Here's one way to understand some of these effects. If you have a model for the magnets-in-coils as a velocity transducer, from section 1.4, you can adapt it to the new geometry you're using here. Then theoretically, and experimentally, you ought to be able to show that small oscillations of the rotor about its equilibrium position, at 'natural frequency', ω_0, will induce emfs in the coil which are of frequency, $2\omega_0$. If you were to connect an external resistor to the coils, there would also be currents of frequency, $2\omega_0$, in phase with the emf -- and there'd also be dissipation, in the resistor, of the mechanical energy of the oscillator. Now imagine removing the resistor, and using an external generator to drive currents in the coil, still of frequency, $2\omega_0$, but which are 180° *opposite* in phase to those just discussed. Because of this flip in phase, these currents must be of a character so as to *add* energy to the mechanical oscillation. This would be called 'on-resonance parametric drive' of the oscillator -- and note that the right *phase* of the generator is required for it to pump up the energy of the oscillator.

The theory of 'parametric drive' is rather involved, so it might be best to discover all these facts empirically first -- that way, the theory will mean something to you when you wade into its mathematics. So set up the Torsional Oscillator using magnetic damping to achieve a not-vastly-high Q, of about 20 to 40, and use the methods of your choice to find a decent estimate of the Q. It turns out that a fractional tolerance of order 1/Q will tell you how far in frequency from the resonance condition you can afford to be. If, in your geometry, you've measured that 'balancing current', i_b, introduced in section 5.5, it will turn out that the threshold amplitude required of the drive current is of order $(1/Q)\ i_b$. In order to see what's going on, and achieve the relatively fussy conditions required for parametric excitation, you might want to use a real-time XY-display on an oscilloscope, with $V_{pos}(t)$ on the horizontal axis, and the drive waveform (or its surrogate) on the vertical axis. You may also want to hand-excite the oscillator at a modest amplitude, on the order of 0.1 radian, so as to have some signals to see.

Your display will show a certain kind of 'Lissajous figure' which is stable when you've reached the required condition for the frequency of excitation. You will eventually find a condition in which you can be persuaded that the angular-position is growing in amplitude, and growing exponentially too (though probably with a slow growth rate). If you think you've achieved parametric drive, here's a really convincing test that exposes the phase sensitivity of the method. Have a reversing switch in place, so that you can suddenly reverse the connections to the drive coil. Then, when you think you're in the exponential-growth mode, *flip* that switch, so that suddenly you're driving with the same frequency and amplitude, but reversed-in-phase. Your oscillator should now go into exponential *decay* of its amplitude, and at a decay rate *faster* than that due to the damping magnets alone.

In practice, exponential growth in the oscillator can't go on too long, since the interaction between the coil and the magnets on the rotor is only simple in the small-angle approximation. Similarly, exponential decay induced by reversed-phase parametric drive will not go down to zero, but will eventually break into exponential growth of a fresh oscillating mode that *is* in the right phase relationship with the drive. But short of these limiting cases, you should be able to get time records that do show characteristic growth and decay rates, and you should be able to quantify these rates. In fact, since this system lacks anything corresponding to steady-state amplitude, such growth (or decay) rates are your chief observable, the dependent variable you can measure. Your *in*dependent variables are the frequency and amplitude of the drive waveform. You should be able to confirm the claims above:

- as to frequency, you should (at any fixed amplitude above threshold) find there's only a *finite* range around the 'resonant', i.e., 2-to-1, frequency at which you can get growth in oscillation, and this range increases with amplitude, but decreases with the Q of the oscillator you're driving;

- as to amplitude, specializing to the 'resonant', i.e., 2-to-1 condition, you should be able to find the *threshold* amplitude, and show that it is proportional to (1/Q) i_b as claimed above.

Finally, there are details as to the phase shift that will exist between the drive and the response. You might see this empirically first, since it's hard to think theoretically about the meaning of a phase shift between two signals that differ by a factor of two in frequency. But you should confirm that when you are in the exponential-growth mode, the oscillator is running not at its natural frequency, but at half the drive frequency, and that you can detune the drive a bit from your target (double the natural frequency of the oscillator) at the cost of this phase shift. In fact, the oscillator will be phase-locked to half the drive frequency, with a phase error that is zero right at 'resonance', and finite and stable in time when you're a bit away from 'resonance'.

This section has not begun to introduce even the equation of motion that describes your oscillator under these conditions, but rather leaves it to you to work out. You will find that it is *not* a second-order differential equation with constant coefficients, inhomogeneous because it's driven by a function, as formerly. Instead it's a second-order differential equation with *non*-constant coefficients, and homogeneous in character. In fact, you can think of it as the equation of an *un*driven oscillator, except that one of the parameters of the system has become a time-dependent function. (Hence the term 'parametric excitation'.) If you search the appropriate references, you will be amazed at the mathematical depth at which all of this can be understood, using artful approximations in an analytic treatment. Alternatively, you can find the mathematical literature on the Mathieu Equation. You'll also be amazed if you look up the exotic applications of parametric drive, parametric excitation, and parametric oscillation that turn up in curious branches of science and technology.

5.7 Coupled oscillators

There are lots of cases in physics that can be modeled as two simple harmonic oscillators which are more or less *coupled* to each other. The phenomena that emerge from this coupling are of widespread interest in classical and quantum physics. If you have two Torsional Oscillators, you're ready to be able to investigate these phenomena, both qualitatively and quantitatively.

To perform this investigation, you need to have the rotor magnets' moments aligned perpendicularly to the Helmholtz coils' axis, as they are shipped (and as they are used everywhere except in sections 5.4 - 6). You'll need to know how to 'tune' the frequency of oscillation of undamped oscillators, after the fashion of section 1.3, by adding masses to the rotors. You have brass quadrants for coarse adjustment, steel balls for finer adjustments, and the possibility of adding steel balls of smaller diameters for smaller adjustments still. You'll want your two oscillators to have natural frequencies that match to better than 1%, and at this level there are issues that arise, such as the location and orientation of the oscillator -- recall section 2.7.3.

To start this investigation, you want to reduce the damping of both oscillators to a bare minimum, and you'll need to place two Oscillators side-by-side on a tabletop, with their bases nearly touching -- label them and their locations by #1 and #2. You want to match two frequencies: that of oscillator #1 at location #1 (measured when oscillator #2 is distant), and that of oscillator #2 at location #2 (measured when oscillator #1 is distant). The reason for doing the measurements this oddball way is that the oscillators interact when they're adjacent! In fact, before going on, you should work out the magnetic energy of interaction for two dipoles, μ_1 at angular orientation θ_1, and μ_2 at angular orientation θ_2, when their centers are separated by vector, **r**. You can treat the two dipoles as point-like to adequate accuracy, and after getting the exact $U(\theta_1, \theta_2)$ function, you should expand it for small (θ_1, θ_2) values. Now write the total potential energy of the system of two oscillators, and show that it has the form

$$U(\theta_1, \theta_2) = (1/2)\,\kappa_1\,\theta_1^2 + (1/2)\,\kappa_2\,\theta_2^2 + c\,\theta_1\,\theta_2$$

which is just the form of coupling that leads to the simplest possible interaction between two oscillators. Notice that your modeling of the magnetic moments gives the coupling constant, c, as a predictable number, and you can even predict how it should change with the separation of the two oscillators.

Now for some qualitative observations. Embark on a series of observations, each of which starts with oscillator #1's rotor held away from equilibrium, and oscillator #2 hand-damped to quiescence. Now release #1's rotor and observe. You should see oscillator #1 start oscillating, and oscillator #2 start to 'awaken from its sleep'. You will see a slow cycle of behavior that reduces the energy of #1 to a minimum and raises that of #2 to a maximum, followed by a reverse flow of energy back into #1. You can fine-tune the natural frequency of either oscillator until you achieve the desired matching condition: at this condition, the energy of #1 in its slow cycle will not only reach a minimum, but that minimum's value will be *zero*. This is quite a remarkable phenomenon to watch! You might not be surprised to see a more energetic oscillator give energy to another of smaller energy, but in the late stages of the first energy exchange, you're seeing an oscillator with little energy give up *even that little energy it has* to another oscillator of greater energy! It's a bit like seeing water flowing uphill spontaneously.

You can easily measure the 'recurrence time' in this system of coupled oscillators -- that's the time from having all the energy in oscillator #1, until it's all back in #1 again -- and if you work out the theory of coupled oscillators, you'll find that this recurrence time can be predicted in terms of the coupling constant, c. In fact, this method probably offers you the most precise way of measuring the constant.

There are lots of other quantitative observations you can make, including novel kinds of 'phase plane' plots, such as the locus of the $[\theta_1(t), \theta_2(t)]$ points with time. But real progress comes from understanding this system in terms of 'normal modes'. If you have tuned your system to the matching condition discussed above, the normal modes are quite simple and can easily be excited by hand. One of them is the 'symmetric mode', the other the 'anti-symmetric mode', and in both modes, the amplitudes of the two oscillators' motions will be equal in magnitude. What's novel is that the two normal modes have *different frequencies*, and the difference in these frequencies can be predicted from the value of the coupling constant c.

Yet another check of the coupling constant can be made in an entirely static measurement. From the form of the potential-energy function, $U(\theta_1, \theta_2)$, above, you can find the 'global minimum'; but you can also show that if θ_1 is hand-turned to a small static value (such as $\theta_1 = 0.1$ radian), then the system will have lowest energy if θ_2 *also* departs from zero, and by an amount calculable in terms of c. The effect on θ_2's position is not large, but readily detectable using the high sensitivity of the angular-position transducers.

There is another clever trick you can use to understand and even 'purify' the normal modes. As you've set them up thus far, they both have minimal damping. And as you've used the Oscillators thus far, you've had the drive coils either un-connected, or used as velocity transducers. But now bring out the coil connections directly (via the grey-banana terminals) and simply connect one coil to the other, via two wires. In one of the two normal modes, the emf's separately generated in the two coils will be 'pointing' in opposite senses, so the two emf's will cancel each other out, leading to a negligible current flowing in the circuit you have established. In the other normal mode, the two emf's will be 'pointing in the same sense', and will thus drive a non-zero current through the two coils. Clearly, this mode will have a larger rate of energy loss, due to Joule heating. The result is that the two normal modes still exist, but they show up with quite different damping constants! And just by interchanging the two leads connecting the coils, you can interchange which mode is high-loss, and which is low-loss.

Now, you will need less skill to get a pure normal mode. You could excite the system *any way you like*, say by a hold-and-release of only one of the two rotors. That will set up a superposition of the two normal modes, with the energy-exchange consequences you saw earlier. But, if you connect the coils, the energy in one of those two modes will be much more rapidly depleted than the other, thus soon leaving the low-loss mode in 'purified' form. Now you can disconnect the wires, leaving the system in the normal mode you've created.

(If you know the physics of the neutral-kaon system, you can translate this language of excitation and normal-mode loss rates into the vocabulary of K_1 and K_2 mesons, and K_L and K_S eigenstates.)

Thus far you've aimed for, and exploited, a 'tuning' of the system such that the two oscillators, if isolated from each other, would display equal frequencies. But you can also fine-tune the rotational inertia of either oscillator, by small and precisely-calculable amounts, and explore the 'spectroscopy' of the system. For each setting of 'tuning', you can measure the frequency of both normal modes -- by this point, you will have a good method for exciting 'pure versions' of both normal modes in turn. You could label the 'isolated frequencies' of the two oscillators as ω_1 and ω_2, and the frequencies of the two normal modes as ω_a and ω_b. Then you should show theoretically that the combinations

$$\omega_a^2 + \omega_b^2, \text{ and } (\omega_a^2 - \omega_b^2)^2,$$

are predicted to be particularly simple functions of the system's parameters. In fact, if you graph each of these combination as function of 1/I, where I is the rotational inertia of the rotor you are varying, you will be able to extract a set of parameters that describe the system neatly. Finally, you should make a plot of how ω_a^2 and ω_b^2 vary as a function of 1/I, to show a beautiful 'avoided crossing' of the sort that shows up in related coupled-oscillator systems in quantum mechanics.

6 Specifications

Here's a collection of information about parts in the Torsional Oscillator.

6.0 Surfaces and care of the Torsional Oscillator

This section of the manual tells you how the surfaces of the various parts of your Torsional Oscillator are finished, and suggests what care they'll need accordingly.

- the wooden box and base are finished with polyurethane, and will only need dusting; the coating is resistant to moisture, and reasonably hard.
- the aluminum pieces are mostly hard-anodized, some 'clear' and some black in color. This makes them highly resistant to scratching. By exception, the clamps that hold the torsion fiber, and the card guides that hold the edges of the position transducer, are purposely left as uncoated aluminum.
- the torsion fibers are hard-drawn low-carbon steel 'music wire', and are phosphate-treated to leave behind a slight surface film. This film should be left on the fibers, to keep them rust-resistant. Liquid moisture should be kept off the fibers, since they are not rust-proof in wet conditions.
- the copper rotor disc is high-conductivity copper, and it would rapidly tarnish in air, except that it's been 'hard powder coated' and thus ought to retain its color indefinitely. The coating could be scratched with tools, but certainly won't be damaged by touching.
- the brass quadrants and brass weights and hang-downs are unfinished as-machined brass. They will retain their color and finish in indoor environments, but will eventually tarnish under fingerprints. The effect is purely cosmetic, and will scarcely change the mass significantly. Commercial brass polish could be used to restore a shiny finish.
- the steel balls are 'chrome steel', typical of most ball bearings. They are supplied with an oily surface film, and will stay rust-free if this is left on. Removal of the film, and exposure to very moist air, would eventually cause surface rusting of the steel; this can be removed, if desired, with steel wool.
- the permanent magnets are nickel-plated, and ought to remain corrosion-free indefinitely. In the magnetic dampers, the 'magnet arcs' joining the pairs of magnets are steel, also nickel-plated.
- the front panel is brass, with a brushed and protected front surface that should neither corrode nor require maintenance.

Finally, various machine screws apparently sunk into the wooden box are in fact fastening into threaded metal inserts, themselves press-fit into the plywood. Thus such screws could be removed and re-inserted without fear. If such a screw were to be grossly over-tightened, the metal insert might slip relative to the plywood -- so take reasonable care if you should have occasion to tighten such screws.

6.1 Masses and sizes of relevant parts

This section gathers into one place a number of relevant dimensions and masses of various parts of the Torsional Oscillator. Masses are given in grams, and dimensions are quoted (for cultural reasons) in inches, where 1 inch = 1" ≡ 0.0254 m exactly.

- the copper rotor disc has maximal outer diameter of 4.95", an inner diameter of 1.02", and total mass of 962 ± 2 g. The diameter of the circle of holes receiving the brass quadrants' dowel pins is 2.720" ± 0.005".
- the rotating part of the angular-position sensor has outer diameter of 4.74", an inner diameter of 1.02", and total mass of 37 ± 1 g. It's made of standard printed-circuit board material, 1/16"-thick epoxy-reinforced fiberglass (FR4), with copper electrodes.
- the brass quadrants, as supplied with stainless-steel dowel pins in place, and as mounted on the copper rotor disc, describe arcs with an outer diameter of 3.72", an inner diameter of 1.72", and have a total mass of 214.5 ± 0.5 g each.
- the steel balls are chromium-steel bearing balls, of diameter, 1.0000", with amazingly tight tolerances on diameter and roundness. They have mass of 66.8 g each. Placed in the conical depressions atop either the copper rotor disc or the brass quadrants, their centers lie on a circle of diameter, 2.720" ± 0.005".
- the rotor shaft is made of aluminum, and has maximal outer diameter of 1.50", typical outer diameter of 1.00", and a 0.38"-diameter hole through most of the length of its axis. The aluminum part, without magnets or mounting screws, has a mass of approximately 283 g.
- the magnets on the rotor shaft are nickel-plated NdFeB discs, each with a diameter of 1.00" and a thickness of 0.25". The stack of four magnets is separated at its center by a rib, 0.24" thick, that is a part of the rotor shaft. The mass of the four magnets together is 97 ± 1 g.
- the torsion fibers supplied with the apparatus are music wire conforming to ASTM A228, with very tight control over the diameter. The nominal diameters are 0.029", 0.039", 0.047", and 0.055", and these should be reliable in value, and in roundness, to ±0.0005". The fibers are supplied with nominal lengths of 30", and the measured masses of the four fibers are 2.53 g, 4.62 g, 6.56 g, and 9.12 g respectively.
- each 'air paddle' is made of a 20" piece of aluminum tubing, of outer diameter 0.250" and wall thickness of 0.014", and each tube has a measured mass of about 9.2 g. The paddle itself is constructed of foil-covered foam, nominally 6" by 4.5" in size, of a measured mass of 8.4 g. Of the tubing's length, 3" is immersed in the foam.