Preface for the Instructor

This manual describes a host of things that you can do with the equipment included in TeachSpin's "Modern Interferometry" package. The modular optical devices included in this kit are versatile enough that you and your students can use them to build a variety of interferometers, and do a whole series of experiments, but here's a suggestion:

it will rarely be advisable to start at the beginning of this manual,
and merely march through it serially.

Much better would be to appreciate the structure of the manual, with its implications for lab practice. The manual includes "Building" chapters (sections 1, 3, 5, and 7), in which you're led to construct certain optical instruments; it also includes "Interlude" chapters (sections 2, 4, 6, and 8), which aim to teach certain optical and conceptual skills. Sections 9-16 are "Experiments", aimed at displaying some applications of interferometry to a variety of physical measurements. Finally, certain technical information is sequestered in Appendices A-S.

No doubt every novice ought to start with Section 0, for an introduction to interferometry, and then work through Section 1, for gaining hands-on familiarity with the apparatus by building a first interferometer. But the rest of the manual can be used in nearly any order, as desired. Certain of the "Experiments" will motivate the construction of various interferometers discussed in the "Building" sections; certain interests of the students will motivate reading some of the "Interludes". Some suggested sequences of mutually reinforcing topics are listed in Appendix A of this manual. But it's important to realize at the outset that TeachSpin is providing you *not* a fully constructed interferometer, but *rather* the components that will allow construction of several kinds of interferometer, and allow many kinds of measurements. Our view is that hands-on assembly skills, and emerging forms of conceptual understanding, are goals as important for students as 'measurement capability' narrowly understood.

We think it would be ideal for each student to encounter this apparatus by starting with an empty optical breadboard, building up a package of achievements, skills, and forms of understanding that might be unique. There are too many experiments listed in the manual, or latent in the apparatus, to expect any one student to perform them all. So it's perfectly appropriate for an instructor, or student, to be selective, aiming to investigate one sort of interferometer, or measure one kind of physical effect. In our view, the goal shouldn't be to *cover* all of interferometry, but to *uncover* some of the wonderful things that interferometry can do.

Preface for the Student

Perhaps you've read about interferometers in textbooks, but now you're about to encounter interferometry in a hands-on, build-it-yourself laboratory way. Here's a little background about the coverage of this manual and the TeachSpin apparatus it describes.

First point is that there are many kinds of interferometers, and you have the potential to use this equipment to construct at least three distinct kinds of interferometers. And construct them you will; ideally, you'll encounter this equipment starting with an empty optical 'breadboard', on which you'll build the very apparatus you're coming to understand. In the process, you'll gain skills in hands-on optical layout and alignment tasks that are a sort of co-curriculum to these investigations.

Second point is that interferometers can be used to measure many different things, and this manual will introduce you to some of these techniques. You can view an interferometer as a tool for measuring displacements (to the level of micrometers or even nanometers), or for measuring time delays (to the level of femtoseconds or even less), or for measuring optical coherence (and you'll come to understand what that is). In fact there are so many things that various interferometers can measure that it's unlikely you can cover them all; there's no shame in being selective, therefore.

Third point is that this manual aims to be a pretty complete description of the TeachSpin equipment that it covers, but it can't pretend to cover interferometry completely. You'll need to supply your own background reading, references, and theoretical derivations to create compete reports on what you do; in general, you'll need to plan your own forms of data taking too, since this manual describes more of 'what you can set up' than 'how to take data with it'. The exception to this is the derivation, on a few occasions, of the signals that are to be expected from clever uses of the polarization of light; these exercises in 'polarimetric detection' are worked out using complex representations of sinusoids, so as to make the algebra as simple as possible.

Fourth point is that the manual contains a lot of words, and you needn't read them all before you start. Every user should read Section 0, and browse Section 1, before starting in the lab; but much of the manual will work best if you read it for real in the context of hands-on contact with the apparatus, and experience eye-hand-mind coordination in the lab. This is where all the words and abstractions have the chance to come alive, as they describe concrete optical objects and visible optical phenomena that you'll encounter and use.

Interferometry offers a combination perhaps unique to you as a learner: there is a very tight and immediate connection between what you do with your hands and what you see directly with your eyes; and yet you are gaining direct access to the wave character of light, with all the prodigious sensitivity thereby made possible. A good deal of your hands-on experience of set-up and alignment will be conducted with more use of geometrical insight than of algebraic formulae, so if you're good at eye-hand coordination, you're going to enjoy what you're about to encounter.

Modern Interferometry: Nanometers, Femtoseconds, and Coherence

Table of Contents

0 Introduction to Interferometry

Interferometry means 'to measure using interference', and it stands for a whole range of techniques practiced in optics. The TeachSpin apparatus you're about to encounter will teach you some techniques that can be used, and some targets to which they can be applied, in the very versatile technique of measurement using interference of light beams.

The various interferometers that you will build and use all involve a similar physical concept: this is the division of a beam of light into two separated beams, which (after following separate paths in space) are then brought back to recombine. An optical technique like this is called interferometric if the recombination process can be either constructive or destructive; this is where the 'interference' comes from.

The difference between two beams combining constructively, as opposed to destructively, is a matter of phase; you should know that two beams differing by 180 degrees (or π radians) in phase will interfere destructively. You should also realize that to obtain such a phase shift requires a geometrical path difference of half a wavelength, which is (for beams of red light) about 0.3 μm = 0.3×10^{-6} m, or equivalently a travel time difference of half an optical period, which is about 1 femtosecond = 1 fs = 1×10^{-15} s. Numbers like these should illustrate to you how sensitive a measurement technique interferometry can be; what's more, you'll be building interferometers which are sensitive to phase changes a *thousand-fold smaller* than those just mentioned.

You may have first encountered interference in a 'two-slit interference' experiment, and you can recall that its operation does indeed involve two separated beams, or at least two separate identifiable paths of travel, of light. An opaque screen with two or more open slits in it represents a class of interferometers depending on 'division of wavefront', in which the two beams in question come from different places on a single wavefront. The interferometers you're about to encounter depend instead on 'division of amplitude', in which a light beam is (everywhere across its whole wavefront) split into two outgoing beams, each with only a part of the amplitude or intensity of the incoming beam. An essential device in such an interferometer is a beamsplitter, which partially reflects, and partially transmits, a beam of light incident on it.

You will have the opportunity to build interferometers of three distinct topologies, the Michelson, Sagnac, and Mach-Zehnder interferometers. There are yet more interferometric set-ups than these, but these three will show you a range of techniques and applications representative of the whole field. In addition to building these interferometers, you will have the pleasure of learning, on a tabletop and very much in hands-on fashion, a great deal about the arts and crafts of optical layout, assembly, and alignment. Another whole range of skills and insights that you'll encounter are connected to the clever use, in some of these interferometers, of the *polarization* of the beams of light you'll use.

What's so *modern* about 'Modern Interferometry'? After all, Michelson invented his interferometer around 1880, and the other two topologies you'll investigate date from the early 20th century. Here are some of the modern touches in the experiments you'll encounter. First of all, they will all start with laser sources of radiation, and will all involve the electronic detection of the results of optical interference. Another modern feature is that your experiments will involve the quantitative control of independent variables, and the quantitative measurement of dependent variables. Yet another feature will be the use of modern optical components and mounts, such as multilayer dielectric coatings, and kinematic and flexure-based mountings. Finally, you will be introduced to some very clever applications of polarization in light sources, propagation, and detection.

You should know that the applications of interferometry in modern physics and technology are even broader that those which can be illustrated here. Interferometry shows up nowadays in a host of applications involving high-precision use of the fundamental standard of length, or the very sensitive detection of mechanical displacements. Various interferometric methods are used for all sorts of sensitive testing of the shapes of optical surfaces. Interferometry is the basis of the very important method of Fourier-transform spectroscopy, especially useful in the infrared region of the spectrum. Finally, interferometry of superb sensitivity is the basis of the LIGO project, in which truly tiny distortions of space-time itself, produced by the passage of gravitational waves, are to be detected using optical interferometry.

1 The Michelson Interferometer

a. The simplest interferometer

It's very easy to draw the basic topology of a Michelson interferometer, given the concept of a narrow beam of coherent radiation propagating in space, and the idea of a beam splitter. Figure 1-1 below shows a more simplified layout, and deliberately depicts the components involved as somewhat misaligned, just to make the various beams more visible.

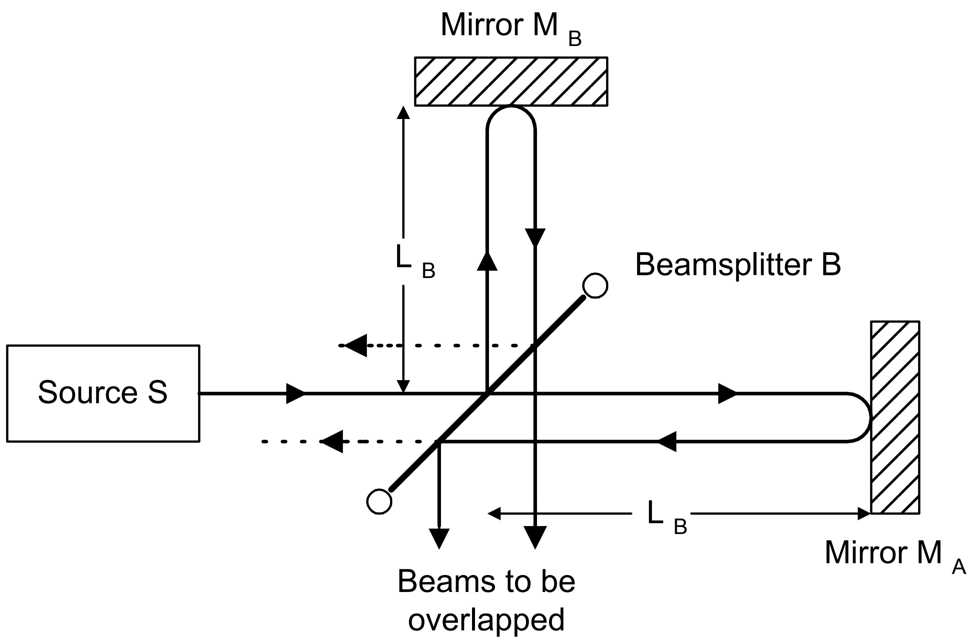

Figure 1-1: A simplified Michelson Interferometer

We start with a source S, producing a monochromatic beam of wavelength λ or frequency f, and some irradiance I_0. We let that beam fall on a beamsplitter plate B at a 45° angle of incidence, so that the partially reflected and partially transmitted beams that emerge are at nearly right angles in space, and are both of irradiance $I_0/2$. Then we let both these beams impinge on fully reflecting mirrors M_A and M_B, so that they are both retro-reflected back toward beamsplitter B. [We consider that the two beams will have traveled one-way distances L_A and L_B (from beamsplitter to mirrors) which may *differ* by an amount that might be a small fraction of a wavelength, or perhaps many meters.] Now at the beamsplitter, each returning beam will again be split into two outgoing beams; these four beams are shown in the diagram. Two of them are headed back toward the source S, and are inconvenient to access; but two others leave the beamsplitter B in a convenient and accessible direction. To *align* an interferometer is to arrange these two outgoing beams to overlap in position and direction. If certain 'coherence requirements' can be met, then these two beams, each of irradiance $I_0/2$, will *not* merely add up in power to yield a beam of irradiance I_0, but instead will superpose, constructively or destructively, to produce a single beam whose irradiance will fall somewhere in the range 0 to I_0. *Where* in this range the irradiance will fall will depend, with interferometric sensitivity, on the value of L_A - L_B.

Now it's time to translate this drawing on paper to a working three-dimensional reality. This requires an introduction to the components available to you in the Modern Interferometry kit, and the technique of handling them safely. You might start by using as light source the Helium-Neon (HeNe) laser that's available to you, mounting it in its holder near a corner of the optical table as shown in Fig. 1-2. If you're the first-ever user of the laser holder, you can assemble it from its parts.

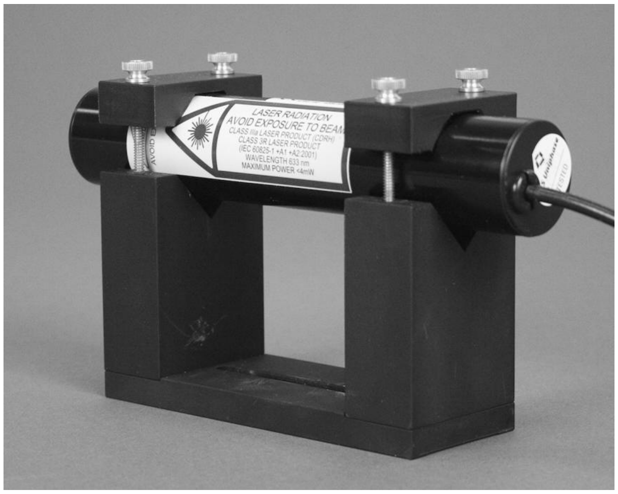

Figure 1-2a: Mounting the HeNe laser

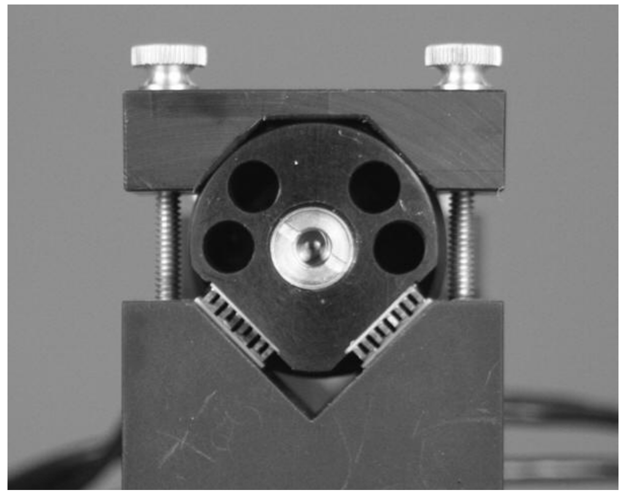
Figure 1-2b: Mounting the Diode Laser

Now locate one of the white-painted 'alignment towers' from the kit, and place it some distance downstream of the laser's output port to form a viewing screen. You'll now need to plug the laser's electrical connector into the mating receptacle at the back of the Modern Interferometry control box, and then find out how to turn on that control box (switch on the back panel) and how to turn on the HeNe laser (toggle switch on front panel). After a few seconds' delay, you should see a fine red beam emerge from the laser. You may now learn how to use the alignment tower as a combination shutter/view screen for the beam propagating across the table.

Now is the time to read and master the safety information in Appendix B, and to understand why you mounted the laser first, and turned it on later. Note that for safety reasons, it's conventional to have laser beams propagate in a horizontal plane, and one that is **not** at typical eye level. It's also conventional to pick a standard for height of beams above the surface of the optical breadboard; in this optical kit, the standard is 3 inches (where 1 inch = 1" is an ancient, non-SI, but perfectly well-defined unit of length, 1 inch ≡ 25.4 mm, much used in optical manufacturing for historical reasons).

This is also a good time for you to become acquainted with the 1/4-20 mounting screws and their socket heads, and the ball-ended socket driver, which are the standard tools for mounting on this sort of optical table. You'll find that there's no need to over-tighten such screws, and they'll provide more than adequate strength and stability using mere hand tightening with the socket-driver. Finally, you might identify those smaller brass thumbnuts which hold the HeNe laser into its mounting bracket, and check that they are snug but not over-tightened.

The simplified diagram above shows the laser beam flying right into the interferometer, but in practice it's necessary to get control over the alignment of this beam. So emulate Fig. 1-3, which shows the two steering mirrors in the kit used in such a way as to deliver a beam that can be fully adjusted, in position and in direction, before it encounters the beamsplitter that marks the entrance to the interferometer proper.

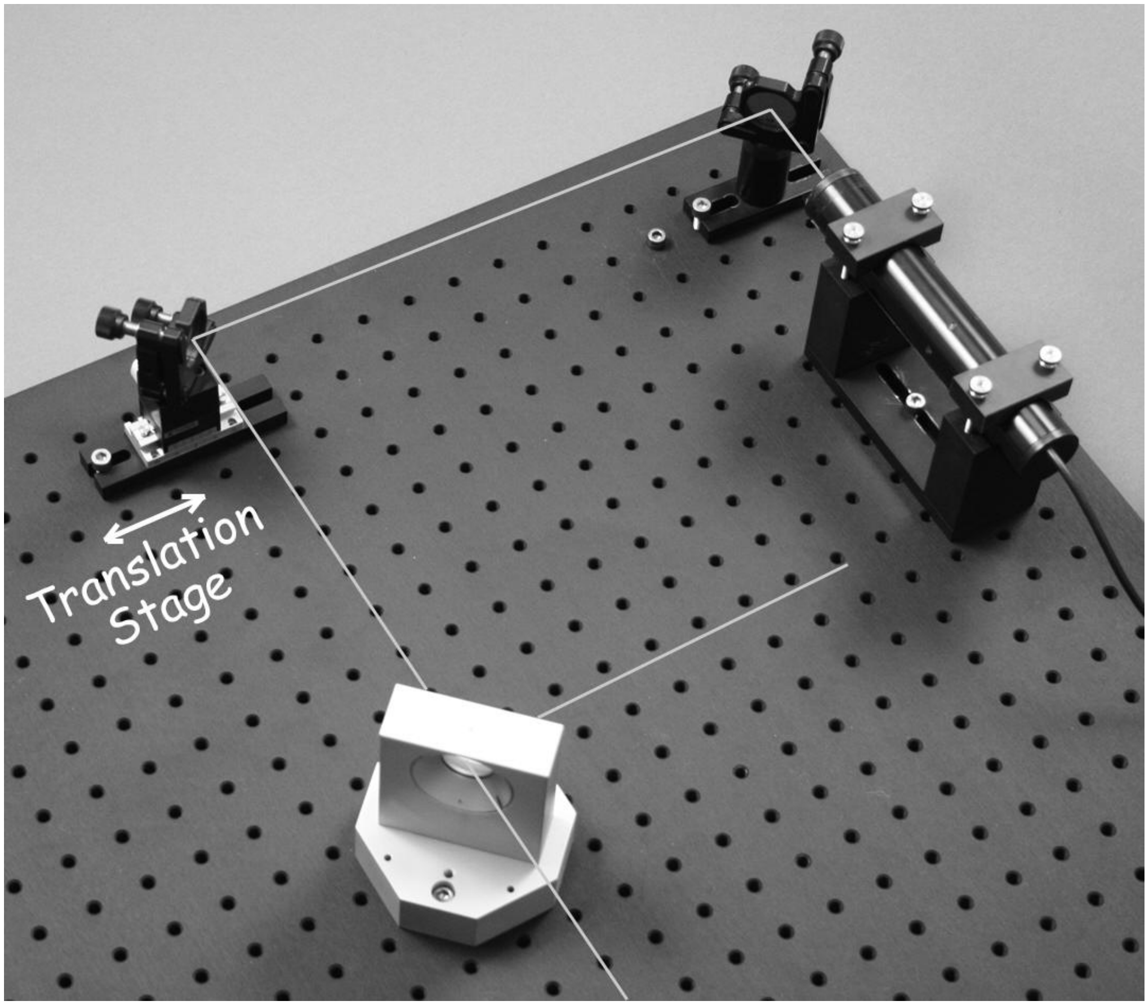

Figure 1-3: One way to use the two steering mirrors

Each of the steering mirrors is a ‘front-surface mirror’, whose (metal) reflective coating is on the *outer* surface of a glass substrate. These metal surfaces, like many others you will see in the Modern Interferometry kit, are working optical surfaces, and are vulnerable to damage either by scratching or contamination, so a basic rule you should hereafter observe is NEVER TO TOUCH AN OPTICAL SURFACE. There are always safe ways to handle optical components that respect this rule; in the present case, you may handle the steering mirrors by their baseplates. You should note that each mirror-mount has two adjustable thumbscrews on its back surface; one of them also has a one-dimensional slide adjustment near its baseplate. For now, you may set these adjustments to mid-range, since you will soon learn how to use optical diagnostics to adjust them.

If you have mounted the beam-steering mirrors to the table properly, you should be able to trace the HeNe laser’s beam through two right-angle bends until it emerges roughly as shown in Fig. 1-3. Now it’s time to mount the beamsplitter, and the two end mirrors, that’ll form the interferometer.

What you want is the base-and-upright structure which holds a thin flat optic of 1" diameter, and which can be used to split the input beam into a transmitted beam and another (reflected) beam, emerging at right angles. (You want to use a dielectric-film 50-50 beamsplitter, not the metal-film beamsplitter you'll use later. See Appendix H on how to tell the difference.) For now, set it down on the optical table in such a way that your beam encounters its optical surface, and use your alignment tower to find the two output beams. Then position the beamsplitter's base above some mounting holes in the table, and screw it down into place. You may have to reposition, or adjust, the steering mirrors to restore the beam to hitting the beamsplitter.

Once you have two output beams, of comparable intensity, you need to choose two end mirrors. From those available in the kit, you need to pick two different kinds, differing in the kind of hinge that they have in their one degree-of-freedom of adjustability. If you can find their adjusting thumbscrew on the back face of the upright, you can see where it bears on the plate holding the mirror, and you can identify the 'flexure hinge' which gives them a combination of extreme rigidity in all other motions, and some flexibility in one rotation. You're looking for one mirror to have a 'horizontal hinge line' and the other to have a 'vertical hinge line', so that your interferometer will have just enough degrees of freedom to allow alignment.

You might put those mirrors down on the table after the fashion of Figure 1-4, and arrange them to retroreflect the laser beams back toward the beamsplitter. Before you attach them to the tabletop with screws, ensure that they define two arms of the interferometer of equal length (to the 1" quantization interval defined by the hole pattern in the table). You might also check that the laser beams are incident near the centers of the front-surface mirrors.

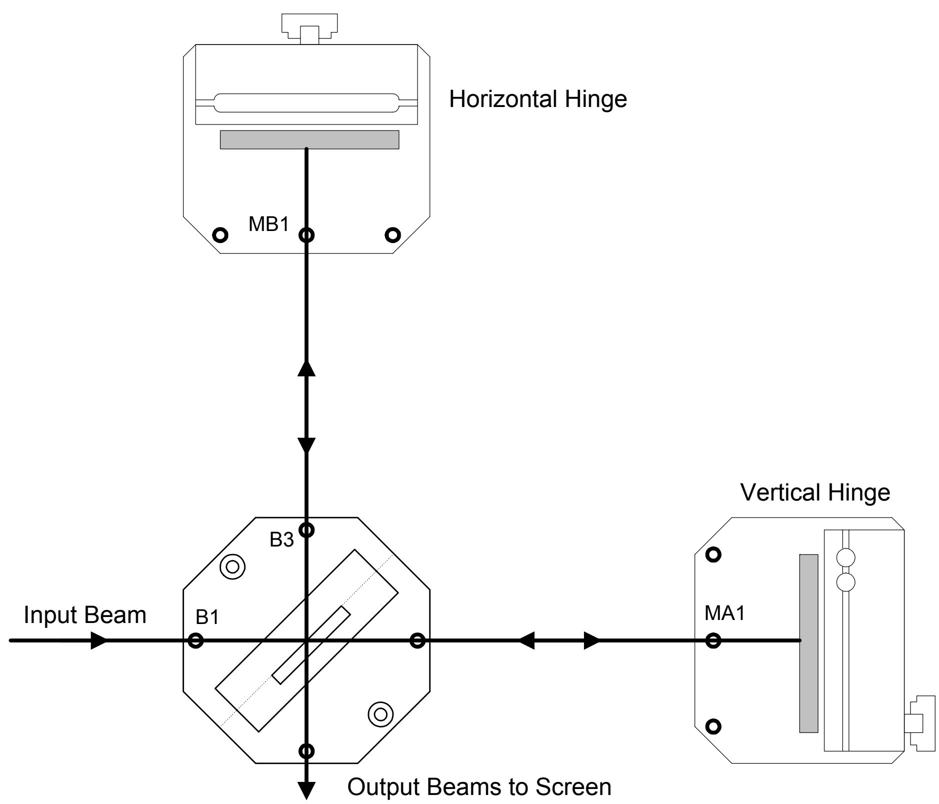

Figure 1-4: A first Michelson interferometer in action

You now have all the components of your interferometer in place, but it is certainly not yet aligned. To do so, you will use not mechanical but *optical* diagnostics, using the laser beam itself as a tool for alignment. The procedure you're about to follow will seem tedious or mysterious only on first encounter, and will later become intuitive and nearly automatic.

1) First you'll want to align the input beam along a natural axis; in this case, it's indicated by the 3-mm-diameter pinholes in two more white-painted alignment tools, called the 'optical paddles'. These have metal dowel-pins on their bottom surfaces, and they slip neatly into the small holes on the tops of the bases of the beamsplitters and end-mirrors, so as to provide alignment locations in space. Put two paddles into place at locations B_1 and M_{A1} as shown in Fig. 1-4, not worrying for now whether the laser beam passes through their pinholes. Now read enough of Section 2 of this manual to understand the algorithm you're going to use: the adjustments are the two thumbscrews on the back of each of the two steering mirrors, and the goal is to get the laser beam passing neatly through the center of *both* of the two pinholes you have established with the paddles. You have aligned a beam of light to be at a standard height above the breadboard, parallel to its surface and along the line of, and vertically above, one of its rows of screw holes.

2) Now reposition the beam paddles into holes labeled B_3 and M_{B1} in Fig. 1-4, to look at the propagation of the beam reflected from the beamsplitter. You'd like this beam also to pass through the central pinholes in the two paddles, yet you seem to have no adjustments left to arrange for it. But you can loosen the screws holding the beamsplitter baseplate to the table, and rotate the beamsplitter holder just enough to arrange a good adjustment horizontally; you can also shim under one corner of the base of the beamsplitter to arrange an adjustment vertically. You may consult Appendix G for practice in doing this better, but it doesn't have to be a perfect adjustment at this stage. When you're satisfied, clamp the beamsplitter mount back down to the table.

3) If your end-mirror mounts are only loosely screwed down to the optical table, you'll be able to rotate them a bit exploiting the 'slop' in their screw holes. While blocking the beam to the other mirror, rotate each mirror mount while looking for its retro-reflected beam to pass back through the beamsplitter's center. You're looking for a moderately coarse adjustment of the mirror-mount as a whole to render its mirror surface perpendicular to the incoming laser beam.

4) If you put an alignment tower into the output beam, you should now be able to identify the (two) standard output beam(s) of the interferometer. At this stage, they are unlikely to be overlapping; this is fine, since this will enable you to identify which beam is due to which arm of the interferometer. [Method: use another alignment tower to block the beam propagating in one arm of the interferometer, and see which spot 'blinks out' when you do so.] Once you have the two beams identified, you can check that each beam can be moved, with one degree of freedom, using the single adjustment screw on the back of the upright of the end-mirror mounts in the interferometer.

5) If you've chosen the right mirror mounts, you'll have one beam spot which can be translated vertically, and a second which can be translated horizontally. Your goal is to adjust each until the two spots overlap. [You have only a finite range of adjustment of the flexure-hinge end-mirror mounts, and if your initial mechanical alignment is poor enough, you might not have enough range of thumbscrew adjustment to achieve this overlap. Rather than damage the flexure hinges by overextending them, you might re-check the earlier stages of mechanical and optical alignment.]

6) When the two spots overlap, you ought to begin to see a new feature, a sort of flashing or blinking of the brightness of the combined spot. This is your first interference phenomenon, and the sign that your interferometer is working. [If the spots overlap with no sign of flickering, move your alignment tower downstream along the beam, and see if your two output beams are perhaps failing to be parallel in direction as well as overlapped in space.]

7) You now have the choice of eyeball or electronic diagnostics for perfecting the alignment. If you're a novice, the eyeball technique is perhaps more educational. You are NOT to look into the laser beams, but you are going to be looking at the millimeter-sized illuminated region on the white screen of the alignment tower for diagnostic information. What you're looking for are 'fringes', or a whole family of parallel stripes alternately bright/red and dark/black. These are formed by the interference of two optical beams whose directions of propagation, and hence whose wavefronts, fail to coincide exactly; they represent interference that is constructive or destructive, depending on position within the light beam. You still have the single adjustment thumbscrews on the back of the two end-mirror mounts as tools, and your goal is to adjust them until the fringe spacing gets as large as possible. When you are done, you'll have no fringes left, just one level of intensity applying to the whole overlap spot of the two beams.

8) You'll know if you're still getting an interference phenomenon, even when the fringe spacing exceeds the spot size, by exploiting the interferometric sensitivity of your optical set-up. Recall that a change in the path difference of only $\lambda/2$ is required to change the output from a state of constructive, to one of destructive, interference. Note that a translation of $\lambda/4$ of one of the end mirrors is enough to yield this change. Now put gentle fingertip pressure on the corner of the metal at the <u>back</u> of the upright of one of the end mirrors, and you'll readily be able to distort it so as to move the mirror the tiny distance ($\approx 0.16\,\mu$m, or about 6 *micro*-inches) required. Alternatively, use some fingertip pressure downward on the optical breadboard itself, to distort it enough to give a similar distance change.

9) Once you've seen the fringes, and used them to optimize alignment, it is time to try out electronic detection of the output beam. Find one of the TeachSpin photodetector assemblies, and see how its mounting post will fit into a post-holder so as to hold its photosensitive surface at the right (≈3") height above the table to capture the output beam. You should mount the photodetector in such a place that you can see the laser-beam spot on the detector's active surface; and if you use the post-holder's baseplate properly, you'll have available both a vertical and horizontal adjustment to make it possible to center the detector on the beam. [You need *not* move the beam to center it on the detector!] Now bring power

to the detector's electronics by plugging its 3-pin connector into one of the three power points available on the front panel of the Modern Interferometry controller unit, and send its signal output via the BNC equipped cable to an oscilloscope. The negative-going output voltage here is proportional to the optical power incident on the photodetector.

10) You'll need to pick a gain setting on the selector switch on the detector, appropriate to the optical power in your output beam. The best diagnostic is your 'scope's view of the signal, which (at a scan rate of perhaps 0.1 s/division) should be showing an indication of the presence of some signal. Use a beam-blocker in the laser beam to establish where the 'zero level' is; then find a way to make the signal vary [as in point 8) above] to see where the maximum is. If your signal reaches -10 V or so, it is saturating the photodetector electronics, and you should reduce the gain setting. You've succeeded when you can see a lively, and very vibration-sensitive, signal on the 'scope, going through a voltage range that extends from near-zero to some maximally negative value.

11) Another alignment technique available to you at this point exploits the vibration sensitivity of the interferometer, and the fast response of the detector. If you have a digital 'scope, you might use multiple sweeps at perhaps 1 s/division to show a time history of the output signal, and you might watch this as you excite the interferometer by repeated fingertip taps onto the optical table. Each individual tap should take the signal through many full cycles of near-sinusoidal variation, as the interferometer distorts by many optical wavelengths under these blows. Your goal is to maximize the contrast of this signal, getting its minima as close to zero, and its maxima as far from zero, as possible. Your independent variables are still the adjustment screws on the back of the two end-mirror uprights. You can at any time use an alignment tower to block the laser beam inside either arm of the interferometer to establish the signal due to one beam alone; if you see a signal level of S_A due to one beam alone, and S_B due to the other beam alone, you should ideally find $S_A \approx S_B = S$, and yet when both beams are present, you should see signals varying over the full range of $(\sqrt{S_A} - \sqrt{S_B})^2$ to $(\sqrt{S_A} + \sqrt{S_B})^2$, or approximately from 0 to 4S.

b. Controlling the fringes

Now to show yourself how far you've come, loosen both end-mirrors' baseplates fully from the optical table, and move both mirrors so as to lengthen each arm of the interferometer by 1". You should be able to tighten down the mirrors and realign the interferometer in a tiny fraction of the time you expended on your first try. If you can demonstrate this facility, you might want to check Section 11 of this manual, *Measuring Indices of Refraction*, to see how you can put an optical element into one beam of your Michelson interferometer, and (finally) thereby introduce controlled and accurately-variable phase changes into that arm, and thereby scan through elegant sinusoidal 'fringes' in your electronic output signal.

Once you have a method (either with rotating a thin glass slab, or using a gas cell of adjustable pressure) for setting the relative phase of the interferometer to a selected value, you can also look to see how stable its output is as a function of time. You will at first be appalled at how unstable your output is, but you will rapidly come to realize the importance of

a) vibration control -- are you working on a vibration-isolated table, of greater or lesser sophistication? See Appendix D for the installation of TeachSpin's simple but effective anti-vibration stiffening structure.

b) air-density control -- do you have masses of air of varying temperature and density wafting through your interferometer? Sure you do, and you may now want to add the draft shield to your interferometer. One corner of this plastic enclosure has holes in adjacent side faces, located so that the input beam from the laser and beam-steering mirrors can enter via one, and the output beam to the photodetector can exit from the other. Note that it's the air in the two arms, not the air in the input or output beams, that needs to be stabilized.

c) temperature control -- have you recently been putting warm hands onto cool metal? Yes you have, and the metal will now be equilibrating back to room temperature, meanwhile changing its dimensions due to thermal contraction. This is a very real and large effect, especially for the aluminum optical breadboard. If one arm (say of 6" = 150 mm length) is mounted on aluminum which rises in temperature by even 1 °C, then it expands in length by about 2×10^{-5}/K, for a one-way length change in one arm of (20×10^{-6}) (150 mm) = 3000 nm = 3 µm, which gives 6 µm change in the two-way travel distance. This is enough to give over 9 full fringes, or 9 full cycles of output variation at your detector. Waiting for equilibration, and the use of the draft cover, will both help to minimize this effect.

With some attention to vibration, air-density, and temperature, you might find that you can achieve stability in your signal at the level of 0.1 fringe over many seconds at a time. This will suffice for many of the experiments you'll do in Sections 9-16; if you are in search of much greater stability, you'll be eager to investigate Section 5, on the Sagnac interferometer.

c. Introducing 'coherence length'

You have thus far built Michelson interferometers with (nominally) equal arm length, and in Section 16 you will learn what can be done with interferometers of *exactly* equal arm length. But for now, there are things to be learnt from building an interferometer of quite unequal arm length. So relative to your equal-arm condition, move one mirror to set up and align an interferometer with arms differing in (one-way) length by 1", and repeat for arms differing in (one-way) length by 3". You may note the markedly smaller fringe contrast for the latter interferometer, and this has a great deal to do with the limited temporal coherence of your HeNe laser.

To see this effect more dramatically, revert to an equal-arm Michelson and line it up, and now substitute (for the HeNe laser you've been using) the diode-laser source that is contained in a short cylindrical aluminum holder. To do this, turn off the HeNe laser (using the front-panel switch) and loosen the four retaining nuts that clamp it into its mount. Slide out the HeNe laser head, and substitute the diode-laser module, clamping it gently under one of the two plastic clamps in the source holder. [See Figure 1-2b in an earlier section.] Now connect its power cord to the Power Output connector on the back panel of the Modern Interferometry controller, and turn the front-panel switch to the Diode position. You should be rewarded with another red beam of light.

You'll note the diode laser output is brighter than the HeNe (nominally 5 mW rather than 1 mW power) and redder than the HeNe (nominally 650 nm rather than 633 nm wavelength). The new beam is also markedly less cylindrical in character, and it will require re-setting the alignment of the two steering mirrors to get it into your interferometer -- use the two beam paddles as before, and your previous algorithm. With modest effort you ought to be able to get a fringe pattern, or fringe signals, from your newly illuminated interferometer. Some details of the physical properties of your two laser sources may be found in Appendix D; in particular, you have some choice for *focus* control of the diode-laser beam.

Now for some issues of coherence length: precisely because the diode laser is quite a bit less monochromatic than the HeNe laser, its output is less coherent temporally, and consequently it matters more that interferometer arm lengths should be nearly equal. If you try (one-way) arm-length differences of 1" and 3" again, you will see the marked difference in fringe contrast, compared with your parallel efforts using the HeNe laser. Section 16 takes up this matter of coherence length in quantitative detail; for now, it's worth thinking about interference as the superposition of two waves which have traveled different distances, and which have thus undergone different time delays. If you have an interferometer of one-way path difference of 1" or 25 mm, then the light reaching the detector is a superposition of two beams, one of which has been delayed by time $\Delta T = 2\ \Delta L/c = 0.05\ \text{m}/(3 \times 10^8\ \text{m/s}) = (1/6)$ ns. [A sixth of a nanosecond may sound like a short time to you, but how many optical periods is that, for red light?] So the question is: does the light the laser emitted one-sixth of a nanosecond ago have a fixed phase relationship with the light it's emitting now? If so, you can get a high-contrast fringe pattern from the superposition of two interfering beams; if not, then the fringe pattern will have limited or zero contrast, as phase variations will wash out the fringe contrast.

A fancy terminology for this process is 'measuring the temporal cross-correlation', and it's important enough that the TeachSpin kit includes provision for varying the interferometric path difference continuously (and not just in 1" increments). Find the translation stage bearing the 0-1 inch micrometer adjustment, and understand that it can substitute for the bottom table-mounting plate of either of your interferometer's end mirrors. The goal is to use it to vary one mirror's position continuously over a 1" range, a range that perhaps passes through the equal-arm condition. [Note again that the actual length of either arm is not the issue; it's the *difference* in length of the two arms that you want to control.]

So learn how to dismount an end-mirror from the optical table, and then separate its mirror-bearing upright from its base. Then find the silvery adaptor plate for the top of the translation stage, attach the upright to the plate, attach the translation stage to the optical table, and (last step) re-attach the adapter plate to the translation stage.

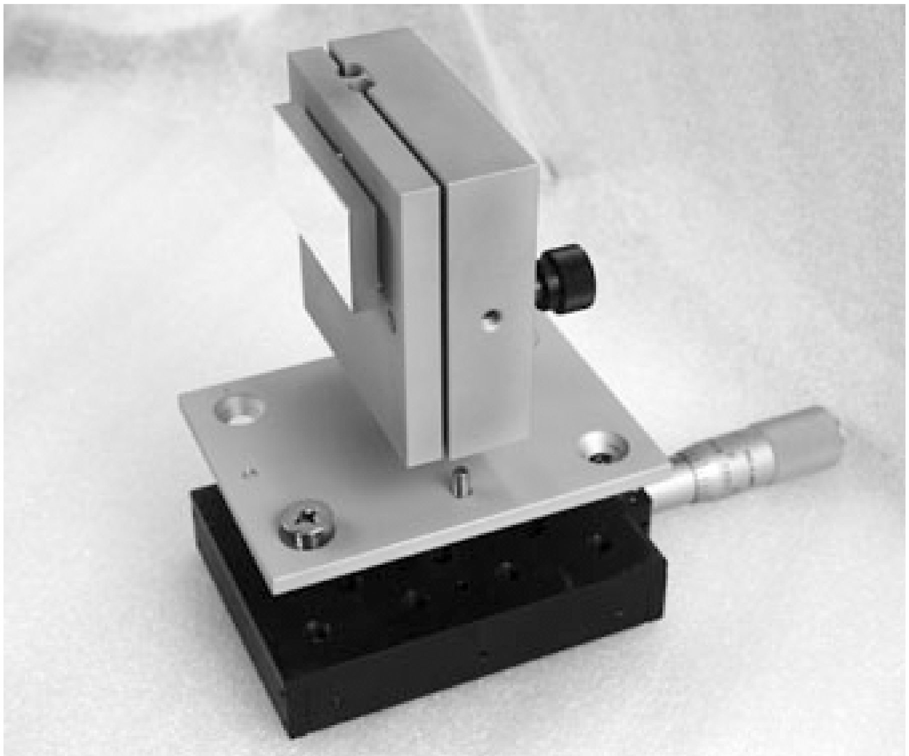

Figure 1-5: Mounting an end-mirror on the 0-1" translation stage

There's a provision for measurement of interferometer arm lengths in the form of a little 'witness mark' on the top surface of your beam-splitter. There a small depression in the metal lies right in the plane of the active surface of the beam splitter. From that point you can measure (to mm accuracy) the distance in space to either end mirror, simply by using a ruler and sighting from above. You might try to be somewhat quantitative about what (one-way) arm-length difference is required to lower the maximum attainable contrast from 100% to (say) 50% for both the HeNe and the diode laser sources. [But since coherence length, and coherence time, are complicated functions of the mode structure of these lasers, and since that mode structure can and does change with time and temperature, it is not worth getting fixated on numerical values here. What *is* important is to see that coherence is a matter of degree, not kind: there is no binary sorting of light sources into coherent vs. incoherent categories.]

If you know enough about HeNe lasers and their longitudinal cavity modes, you can test a very interesting prediction. Your HeNe laser is not perfectly monochromatic, but typically oscillates in two modes with optical frequency near 473,613 GHz, but separated in frequency by 1.090 GHz. As it warms up, it happens that two modes of equal amplitudes and this frequency spacing will be present, with very interesting consequences for interferometry. At certain path-length differences, a Michelson interferometer can display fringes of *zero* contrast, when the maxima of fringes due to laser light of frequency f_1 lie right on top of the minima of fringes due to laser light of frequency f_2. You should be able to show that this first happens at (one-way) arm-length difference near 69 mm, and so you might set up a Michelson interferometer with this target. It will not be impossible to get fringe signals, since (especially while warming up) the laser will not always deliver equal intensity in its two modes. But you should be able to see the fringe contrast pass, with time, through stages of minimal contrast or totally invisible fringes, as the laser passes through the two-equal-modes condition.

If you want to confirm that this is *not* an issue of 'tired light' or some other pathology, you might double the path difference to 138 mm, and explore, in theory and in practice, what this does to the fringe contrast.

If you'd like to see some visually appealing fringe patterns that can be produced with an unequal-arm Michelson interferometer, now is the time to find one of the convex lenses among your collection of equipment, and to introduce it into the laser beam somewhere upstream of the interferometer. The purpose of the lens is to deliver light, downstream in the detection plane, having not plane but spherical wavefronts, and also to deliver light over a much larger area than you've seen illuminated so far -- more like 1 cm^2 than 1 mm^2. You may think of the convex lens you're about to use as focusing the laser beam to a focal point, or more accurately a Gaussian waist, from which it then diverges as desired.

The only challenge is to locate the lens properly; you'll want it coming between the laser and the interferometer's input port, but its location *along* the beam is not crucial. What's a bit harder is to get its position *transverse* to the beam adjusted right, and here's a method. The lens is in a ring mount on a post, and you can put the post into a post holder. Now sliding the post holder across the tabletop will give you horizontal adjustability, while sliding the post up and down in the post-holder will give you vertical adjustability. Your goal is to look at the now-expanded spot or cone of light downstream from the lens, and to position the lens in such a position that this cone is centered on the axis previously defined by your collimated laser beam.

When you've achieved this, you can let that cone of light fall on a alignment tower or other view screen, and see what interference fringes look like now. Practice adjusting the end-mirrors' tilts and the interferometer's path difference, to see what happens as a result. Consult Appendix E if you'd like a mathematical explanation of the characteristic 'bull's-eye' pattern that you can arrange to see.

d. Micrometer-level path-length control

Now for a final episode of path-difference control; for this, you might want to have the HeNe laser as source, and the 0-1" adjustable base in the farther arm of your interferometer. The new goal is to get *really* precise control over the position of the other end mirror, with mechanical resolution in position adjustment of less than 1 μm. You might have noted that the ordinary micrometer driving your present translation stage moves the stage by 0.025", or 635 μm, per full turn, with the result that the smallest division on its barrel corresponds to a (one-way) motion of one twenty-fifth of this, 0.001" or 25.4 μm. Since that gives 50.8 μm of round-trip path difference, or about 80 full fringes for red light, you can see that you don't have very good mechanical control at the single-fringe level. To surpass this level of control, we now introduce you to TeachSpin's monolithic flexure stage, an alternative base for the upright of the (currently stationary) end mirror. The flexure is machined out of a single piece of aluminum, 1" thick, and inside its rectangular outer frame is a central 'stage' or island, which is free to translate by ±1 mm from its central position. [See Appendix F for details on how it's made and why it works.] That stands for only 2 mm (one-way) range, or 4 mm in a round trip, but that is enough to give you many *thousands* of fringes. What's more, the flexure is built in such a way that you can expect to get pure translation of the stage, with unwanted rotations under the 10^{-4} radian level so as to preserve full fringe contrast. Finally, the flexure is intended to be driven by a special differential micrometer or 'diff mike' with 0-2.5 mm range, but readable via a Vernier scale to a resolution of 0.1 μm. This will give you direct mechanical control at the sub-fringe level that your interferometer deserves.

You'll want to read Appendix F on the recommended mechanical set-up of the flexure stage and the 'diff mike', together with the push-spring and the pushrod that keeps the whole assembly in reproducible mechanical contact. Then you'll want to transfer the end-mirror upright from the ordinary base to the flexure stage, and mount the flexure stage back onto the optical table. In order to restrain the flexure stage from free oscillation about its neutral position, you'll also want to engage the push spring, the pushrod, and your micrometer so as to hold the flexure stage near the middle of its range. You might also want to use your 0-1" translation stage, in the other arm of the interferometer, to set the two arm lengths to be nearly equal. Now you should be able to align your interferometer by the usual series of adjustments, coarse and fine.

You'll want to put a detector into the output beam to see the results. You should be able to find a way to exert pure torque on the diff-mike's barrel, and watch the result in the fringe signal. Finally you should be able to control the fringe signal, in the sense of learning how to 'park' the interferometer at a fringe maximum, minimum, or maximal-slope condition. You will also get a feeling for how delicate a touch is required to keep mere fingertip pressure on the micrometer from distorting the optical table in some irrelevant way.

In order that you can deliver a continuous torque to the diff mike's barrel, your TeachSpin kit includes a simple motor-drive unit for this micrometer. It's based on a synchronous motor and gearing unit that delivers a nominally constant 1 rpm shaft rotation, with as little extraneous vibration as can be achieved. The motor unit mates to the differential micrometer via a 'spline drive', which serves two distinct purposes. For one, it accommodates the >20 mm longitudinal translation that the barrel of the diff mike undergoes during its operation; for another, it accommodates the inevitable departure from perfectly coaxial alignment of the diff mike and the motor shaft. One tolerable drawback of the system is that the motor turns in one direction only, in such a way as to back the diff mike away from the translation stage, so as to lengthen that arm of the interferometer. In order that these backing-out motions not lead to a collision between the spline drive and the motor unit, there is a micro switch in place below the motor shaft, whose job is to shut off power to the motor when the micrometer backs out too far. See Figure 1-6 for a view of this motor in position.

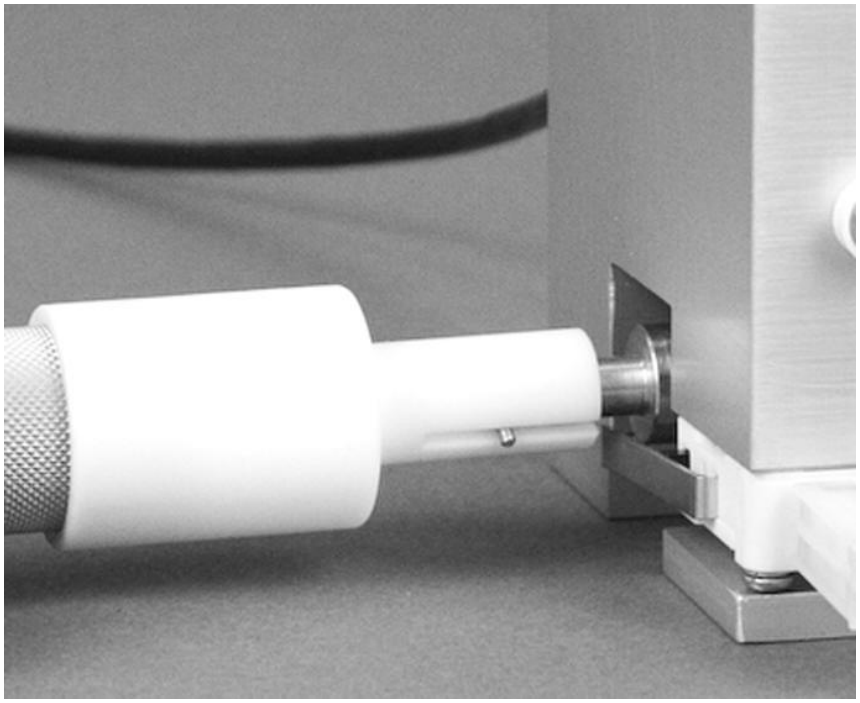

Figure 1-6: The motor drive engaged, via the spline coupling, to the differential micrometer

So you might disengage the whole motor-drive unit, hand-set the micrometer to a desired starting position, and then hand-rotate it to an azimuth such that you can re-engage the motor unit. You'll use the mounting-screw slots in the motor unit to position the motor so its shaft is coaxial with the spline drive, and so the spline drive's drive-pin is (shallowly) engaged in the sleeve on the diff mike; that way, you'll get the full depth of the slot available for translation of the barrel as it backs out under rotation. Finally, you can power up the motor by connecting it to the Motor connector on the back of the Modern Interferometry control box, and actuating the appropriate front-panel switch. Watch to see if you get smooth and steady rotation of the diff mike through a full turn or more; meanwhile, you can watch for smooth and steady sinusoidal fringe signals from the interferometer.

Given the nominal rate of rotation of the motor drive, you can compute the nominal values of the rate of translation of the mirror, the rate of increase of path difference, and the rate at which fringes appear. The motor speed has been chosen such that you can *count* the fringes as they go by, and this allows you to measure optical wavelengths with your micrometer! For best results, you'll want to deal with slack and mechanical backlash by the following procedure:

- choose a starting position for the micrometer, and engage it mechanically as above;
- start the motor drive with*out* counting at first, until the fringe rate has steadied (and the whole mechanical drive train has engaged);
- then use the switch to stop the motor; and while it's stopped, read the differential micrometer's setting;
- now start from a fringe count of zero and restart the motor;
- when you've reached a target number of fringe counts, re-stop the motor, and make the final reading of the micrometer.

You are now accomplishing one of the measurements for which a Michelson interferometer is well suited; for more details on the direct mechanical measurement of the wavelength of light, see Section 9, which takes up this capability in detail.

2 Interlude on *Alignment*

This interlude of the manual has (at first glance) nothing to do with interferometry, but it will introduce you to a whole range of valuable skills in tabletop optics. The basic skill is the task of *optical alignment*, which at its most basic is to ask -- what do I have to do to get a laser beam to shoot right down the axis of a given pipe? You might be thinking about picking up the laser and aiming it, in gun-like fashion, along the pipe, but this is rarely feasible in real-life optics. So to be concrete, go ahead and bolt your HeNe laser source down, somewhere on the table, and then imagine that some other person has used two lumps of modeling clay, and a drinking straw, to lay out a small 'pipe' in an arbitrary but fixed position in the space above your table. How do you get the light beam to pass cleanly down the axis of the straw, if you're not allowed to move either the laser *or* the straw?

The solution typically applied in tabletop optics is to relay the beam from the laser to the pipe using two steering mirrors. Why *two* mirrors, first of all? The answer comes from <u>counting</u>: You have *four* constraints to meet before you've achieved your goal.

> Think of transparent plates, bearing crosshairs, which are glued to both ends of your straw, and now realize that there are x- and y errors, or horizontal and vertical adjustments needed, on the crosshair patterns on *both* ends of the straw. Thus there are four errors that you need to drive to zero, or four constraints you need to satisfy.

Now a typical steering mirror has only *two* adjustment knobs that allow smooth continuous control, so it takes *two* such adjustable mirrors to provide the *four* independent variables you'll need in order to meet the four constraints.

Now it's time for you to translate this guidance into reality, and here's an exercise that will teach your mind, eyes, and hands the procedure. On an otherwise empty optical table, bolt down the HeNe laser at some generic location and orientation, and turn it on. [With low-power, visible-light lasers, we have the luxury of being able to align the system with the beam turned on throughout the procedure.] Now (instead of the drinking-straw task) set the two 'alignment towers' down on the table, at two generic locations perhaps a quarter-meter apart. Think of the line passing through their two alignment holes as defining the 'pipe' along whose axis you want the HeNe laser's beam to pass. [You have the advantage of a 'pipe' with transparent or imaginary walls.] You'll want to rotate the alignment towers about vertical axes until their view screen faces lie face-on to the axis you've established. Now get the two steering mirrors with their bases; they are now your tools for relaying the beam from a fixed laser so that it's aligned with a fixed and given axis.

Of course you have to locate these two mirrors on the optical table, and there are good (and not so good) places to mount them. Here are three suggestions:

1) the light leaving the laser will first encounter the 'upstream' mirror, so it has to be placed where the light beam will intercept its reflecting surface.
2) the light leaving the second, or 'downstream' mirror, has to pass through the 'pipe', so an imaginary line passing <u>backwards</u> through the pipe has to encounter the mirror's reflecting surface. That tells you something about where it needs to be placed.
3) For the sake of the convergence of the algorithm that you'll learn below, it's important to place the downstream mirror *relatively near* the upstream end of the pipe.

So if we label the steering mirrors M_1 and M_2, and the pipe's open ends by E_1 and E_2 (and the light encounters the four of them in that order), the important thing is that the distance M_2-to-E_1 be rather small compared to the distance M_1-to-M_2.

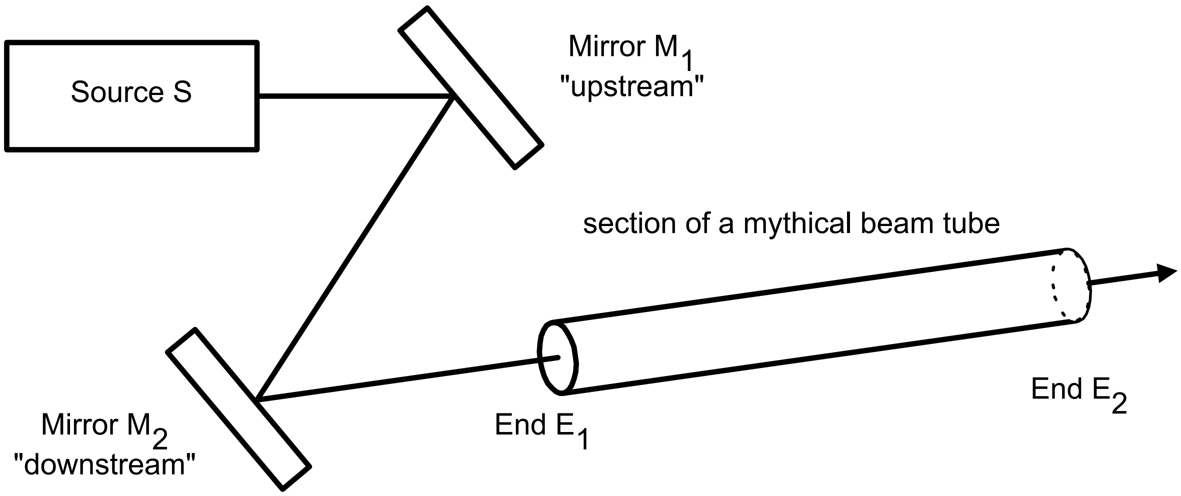

Figure 2-1: One layout for the generic beam alignment task

That tells you where on the table you'll need to bolt down the bases which hold the steering mirrors; now it's time to use the rod holders to adjust the height of these mirrors above the table as needed, and to adjust (by coarse rotational adjustment of their bases or posts) the orientation of the two mirrors as needed. You want light from the laser to hit M_1's surface, and M_1 to be oriented so that light heads toward M_2; then you can use M_1's fine-adjust screws to put the laser beam right onto M_2's surface. Similarly, you can now coarse-adjust the orientation of M_2's surface so that light heads toward the pipe's upstream end E_1 (on the surface of the first alignment tower).

When you have light hitting the upstream alignment tower E_1, use the *upstream* mirror's (M_1's) adjustment screws to get the beam to pass through the hole in the *upstream* alignment tower (E_1). Once the beam is passing through that hole, use a paper card to follow it downstream to the vicinity of the downstream alignment tower (E_2). It may be missing the tower altogether, or hitting the tower's view-screen but missing the hole. Whether by coarse or fine adjustment of the *downstream* mirror M_2, swing the beam around so as to improve its aim toward the *downstream* end of your 'pipe', the alignment tower E_2. It may very well happen that as you adjust M_2, the beam swings properly toward your target E_2, but then gets cut off as it is intercepted by the upstream alignment tower. No matter; continue the adjustment by extrapolation and eye-hand coordination, even if the beam gets cut off. Then go back to the *upstream* mirror M_1 to get the beam in position to pass through the *upstream* end of your 'pipe' at E_1.

You've embarked on an algorithm that will converge rather rapidly (provided that in the course of your fine adjustments, the laser beam doesn't get 'walked off' the edges of a mirror's surface). As you get better at the process, you'll get more experienced in positioning the mirrors such that the beam reflects from points near the center of their surfaces. Remember the key ingredient in the algorithm: you correct errors (in two transverse dimensions) at the *up*stream end of your pipe, E_1, by making adjustments (with two fine-adjust thumbscrews) on the *up*stream mirror, M_1. Then you correct errors (in two transverse dimensions) at the *down*stream end of your pipe, E_2, by making adjustments (with two fine-adjust thumbscrews) on the *down*stream mirror, M_2. Repeat this iterative scheme for several rounds, and you'll see it converge so that it's possible to get the beam to pass right through the center of the holes in both alignment towers.

You'll note that in this exercise you've had the advantage that the laser beam emerges, and the 'pipe' is extended, in the standard optical plane lying 3" above the breadboard. But to see that the procedure works under conditions of greater generality, now reposition the two alignment towers, and this time put some flat objects (of thickness one or two cm) under one or both of them. Now you have a 'pipe' lying at a more random angular orientation, but you'll find the same logic about positioning the mirrors, and the same algorithm for adjusting them coarsely, and then finely and iteratively, will allow you to get the beam to pass along your newly chosen axis.

In actual optical practice in the laboratory, it is convenient and conventional to lay out optical components not only in a layout that keeps laser beams at a fixed and standard nominal height above the table, but also in such a way that beams are directed to lie more or less nearly parallel to the lines of screw holes in its upper surface. This allows the steering mirrors to be used with 45° angles of incidence, so that each mirror deflects the beam through an angle near 90°. Now you can see why in Section 1, you were advised to lay out your first Michelson interferometer with its beams in both arms lying right along the direction, and right above the lines, of the screw holes in the optical breadboard. There are occasions on which there's merit in straying from this 'all right angles' guideline, and if you do so, you will find it much better to have the lasers beams arriving at, and departing from, a mirror's surface forming an *acute* angle rather than an obtuse angle. If you want to try this out, set yourself an alignment task (using the two alignment towers again) in which the beam takes on a Z-folded shape between the laser and the 'pipe' along which you're aligning it.

Now that you have a laser beam neatly aligned with some well-defined chosen axis, there is one more (and much simpler) alignment task you can learn. Suppose you want one of the photodetectors to intercept the beam you have emerging from your 'pipe'. The goal is to align the photosensitive surface to the beam (*not* to move the beam to accommodate the photodetector). To do this, you need to be able to translate the photodetector in horizontal and vertical displacements, preferably while the laser beam on its active surface is visible to your eye. How will you make the necessary adjustments? Given the rather large size of the sensitive area on the photodetectors you're using, it's not necessary to have fine-adjust knobs to accomplish this alignment task. Instead, you can use the rod-in-post holder clamping screw to raise or lower the detector to the height required, and then you can slide the detector assembly, base and all, along the table surface to make the lateral adjustment required. If you align the slotted holes in the post holder base so that their long axes are lying perpendicular to the laser beam's direction, then you'll have a way to slide the detector assembly transversely, as required, while its bolt-down screws are loosened. Remember that you are free to move the whole detector upstream or downstream along the laser beam by the <1" required to gain access to a set of screw holes in the breadboard. [This whole process is easier to accomplish by hand than it is to write, or read, about in words -- try it out.]

There's another scheme related to the alignment of photodetectors, in dealing with ambient room light that can add unwanted background to your laser-beam signals. Precisely because desired laser light reaches the detectors' photosensitive surfaces in the form of a collimated beam, while room light comes in from all angles, it's feasible to pass all the laser light while blocking a good fraction of the ambient light. The TeachSpin photodetectors come with black plastic 'snouts' in place; these slide into (or out of) the aluminum housings of the detectors. The snouts are best removed for initial alignment of the laser beams onto the photodetectors, but they can then be replaced for rejecting some of the ambient light. Naturally, you can build custom snouts out of black paper or other tubing for even better room-light rejection.

Finally, you may want to consult two Appendices relevant to optical alignment. Appendix G discusses the meaning and usefulness of 'shimming' in the optical context, and it will teach you about gaining access to degrees of freedom of adjustment that your mirror mounts (and beamsplitters) seem to lack. Appendix L describes the curious mounting holes that you'll find in the baseplates of your mirror mounts, and how they allow the angular orientations of the mirrors to be set or varied.

3 The 'quadrature Michelson' interferometer

This section assumes that you've built and used a Michelson interferometer, and now poses to you some conceptual and practical questions:

> Q1: When you are at a 'fringe minimum' in operating a Michelson interferometer, ideally there's no energy emerging toward the detector -- so where *is* the energy of the input light going?
>
> Q2: Given the sinusoidal signal emerging from a detector when one mirror of a Michelson interferometer is being scanned smoothly, how can you tell which *direction* the interferometer is scanning?

The answers to these questions motivate a closer look at a Michelson interferometer, and the payoff if a markedly more useful instrument. In particular, you'll be able to build an interferometer so robust against vibration that you can bounce a tennis ball onto its optical table during operation, meanwhile maintaining position resolution under 0.1 μm!

a. Standard, and Non-standard, outputs of a Michelson interferometer

Here's the key to answering those questions, and gaining that practical payoff -- you need to think about the 'other output' of a Michelson interferometer. If you sketch beam paths for such an interferometer, you note two beams, making their turnarounds at the end mirrors, each headed back toward the beamsplitter. You've been used to following, at this second encounter with the beamsplitter, those partial beams that are headed toward the detector; now it's time to follow those partial beams, which are headed the *other* way. Show in your sketch that they are headed back to the laser source in your interferometer; further show that to align the interferometer is to get those two beams headed back to the laser to overlap in direction. Now it's time to study those beams experimentally.

You might best set up an equal-arm Michelson interferometer as in Section 1, using the HeNe laser as a source, but this time leaving 6" or more of clear table space between the second steering mirror and the beamsplitter. Align this interferometer as usual, and now have a look at the black circular output face of the HeNe laser, seeing if you can identify on that face some 'return beams' from the interferometer. Especially if you tilt the end mirrors away from the aligned condition, you should be able to identify the two independent spots arising from beams that are 'folded' at the two end mirrors. Note what happens to these two spots as you align the interferometer (by the standard procedure); you should see them overlap and even interfere. If your alignment is particularly lucky, you may see both beams disappear right into the laser's output aperture, but it's convenient for these observations if they remain visible. You might be able to see the first hint of an answer to Q1 above by asking yourself: when the standard output beam reaches its maximum intensity, is the overlapped-beam spot on the laser's output face at a minimum or a maximum in intensity? And why should this be so?

To look at this effect quantitatively, it's necessary to sample the 'return beam' headed back to the laser and get it to a detector. This is hard to do in a 100% efficient way, since it overlaps in space with the main input beam; but it's very easy to get access to the *relative* power in this return beam, at the cost of some photon efficiency. The method is to use an extra, partially-reflecting, mirror placed at a 45° angle of incidence in the input beam, as shown in Figure 3-1 below. Of course the input beam encounters this mirror on its way *to* the interferometer, and (nominally) 50% of the input light is immediately shunted aside, to miss the interferometer entirely. This is a waste, but the payoff is at hand. The return beam headed back to the laser *also* has 50% of its power deflected by this mirror, and deflected into a place where it's easy to sample with another photodetector. So find a second photodetector, power it up like the other one monitoring the 'standard output' of your interferometer, and use a dual-trace oscilloscope to monitor both outputs simultaneously. Now use the steering mirrors to line up the input beam into the interferometer as before, and use the end mirrors to align the interferometer; you should now be able to see fringe signals at *both* detectors. You should be able to confirm that the two signals are 'complementary', in the sense that when the standard output is at its minimum, the new non-standard output beam is at its maximum.

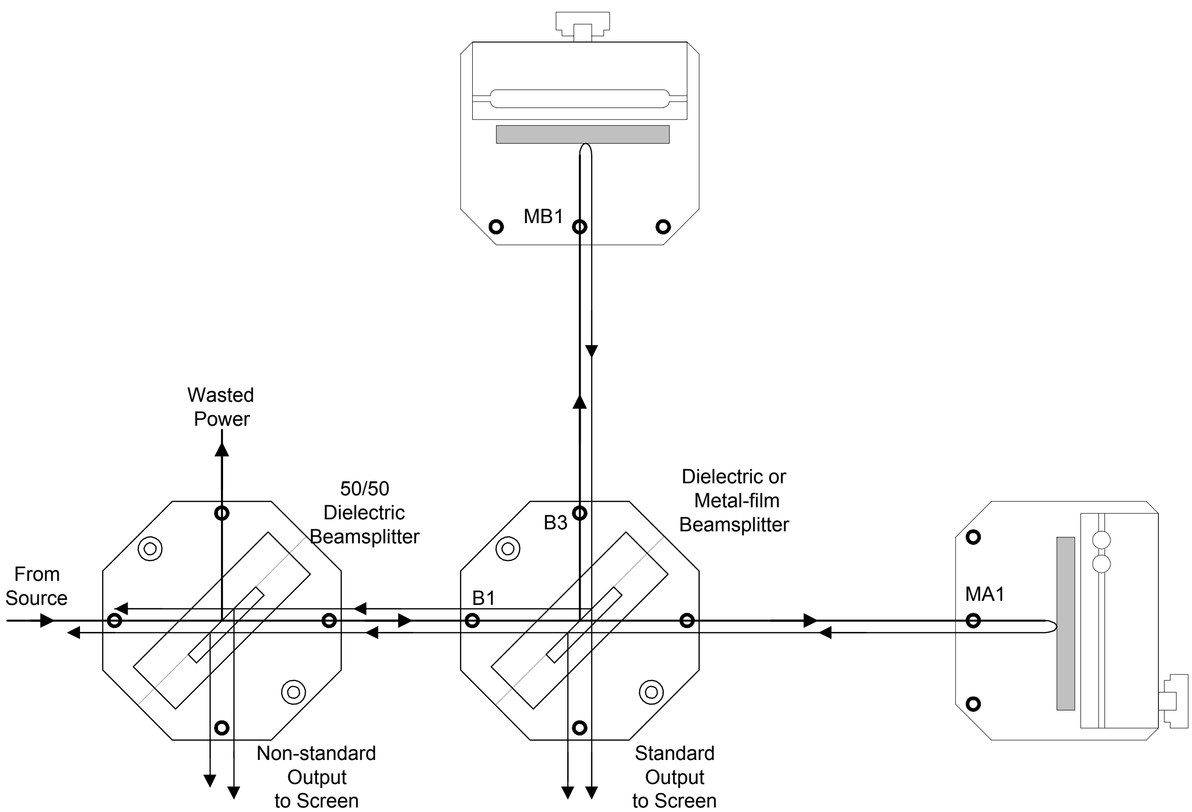

Figure 3-1 Sampling the non-standard output of a Michelson interferometer

If you have these two signals live on a dual-trace 'scope, you should change the 'scope display to the XY mode, in which you can see both signals at once. What is your expectation for these signals? Along what locus ought the instantaneous (X,Y) point to lie? Can you confirm that it does so? Do you see that from the information-theory point of view, that one of the two signals is therefore redundant?

b. Using a metal-film beamsplitter

This is all a lead-in to the new capabilities you are about to explore. The only change you need to make is to take out the (nominally 50-50) multilayer dielectric beamsplitter plate at the entrance to your interferometer, and substitute for it a glass plate with a partially reflecting metal-film coating. [You'll want to leave in place the upstream 50-50 dielectric beamsplitter that is sampling the non-standard output for you.] See Appendix H for information on how to remove and replace these 2-mm-thick beamsplitter plates from their mounting in the beamsplitter upright; the goal is to handle them safely.

Once you have the new metal-film beamsplitter in place, you will need to be aware of the beam reflected from the metal face (as opposed to the residual beam reflected from the glass/air interface on the other side of the thin plate). Line up your input beam with the interferometer using two beam paddles as usual, and then line up your end mirrors until you see the fringe pattern -- look at the *non-*standard output, where the contrast is higher. Adjust as usual for 'zero spatial frequency' in the fringe pattern, and now ensure that you can get signals from the photodetectors monitoring the standard, and the non-standard, outputs of the interferometer. A dual-trace view of the two signals can show you that both go through sinusoidal fringes as you systematically shift the phase difference in the interferometer, but you might notice that the *contrast* of these signals is lower than formerly. That is to say, the minima of these signals do not lie close to zero, but higher. [It's conventional to denote the maximal and minimal values of a photodetector signal's sinusoidal variation as S_{max} and S_{min}, and to define the visibility of the fringe signal by $V = (S_{max} - S_{min})/(S_{max} + S_{min})$; experimentally, it's easier to measure the contrast, $C = S_{max}/S_{min}$. Of course the visibility comes out to $V = (C-1)/(C+1)$. You might formerly have had V nearly 1, but now it may be quite a bit smaller.] Once you can see these fringe signals, it's well worth checking how they behave with respect to your choice of polarization of the input light; be sure to try vertical and horizontal choices of polarization direction, and find the choice, and the alignment, that gives useful contrast for both signals. [More coverage of the polarization of your laser beams' light is found in Section 4 of this manual.]

Now change your display from X-and-Y vs. time to a real-time X vs. Y plot, and you'll see something new, the whole motivation for, and payoff of, the metal film beamsplitter. You should see a new locus for X vs. Y, one that is *not* constrained by a conservation-of-energy argument to lie along a line. That's because (unlike the lossless all-dielectric beamsplitter you were using) the metal-film beamsplitter is dissipative -- you can confirm this by a separate quantitative measurement of optical power going into, and power in the two beams coming out of, the beamsplitter. And free of the conservation-of-energy constraint, the X vs. Y signal can now take on a shape in which the two signals are *not* redundant; so knowing one of them no longer fixes the value of the other.

Seeing what you now see, you can imagine that 'opening up' your locus to the maximum extent possible is the new goal of your alignment procedure. You might tap semi-continually on the breadboard to provide the vibration that will cause the instantaneous (X,Y) point dance around its locus; in the process you might see that the point settles back to its original steady-state location after a tap to the breadboard. Next you will want to make a long slow change of the optical phase difference in the interferometer, perhaps by using the gas cell, or the flexure translation stage, to see how the (X,Y) point responds. Finally you'll have an answer to question Q2 at the beginning of this section, and an extremely valuable interferometric capability.

c. Manipulating and counting quadrature signals

The signals you are seeing are called 'quadrature' signals, since they are mutually out of phase, not by 180° but by about 90°. [A phase shift near 90° is obtained by the right choice of metal film, its thickness, and light polarization.] You can find some of the mathematics of near-quadrature signals in Appendix J, but watching your (X,Y) point move smoothly and reversibly around your locus on the 'scope should give you some really good ideas immediately.

Send those detector signals into the Modern Interferometry controller box, and learn how to adjust them with a variable offset until your locus is centered around the (0,0) point -- again, an X vs. Y 'scope display of the offset-corrected signals is your best diagnostic. You might also adjust the gain until the two signals are of comparable amplitude. Now watch the new point in an X vs. Y display, and ask yourself -- what do you have to do to the interferometer to get one full tracing of the perimeter of the locus, one full 'cycle'? Notice that you have, during one such 'cycle', four distinct events, as the point (X,Y) crosses the axes of your Cartesian display.

These axis-crossings are zero-crossings of electronic signals, and your controller box is equipped to notice them. Send the offset- and scale-adjusted signals you've been monitoring onward to the 'quadrature counter' section of your controller, set the counter to the quadrature (rather than channel-A) mode, and see what the 5-digit counter display does, each time the point (X,Y) crosses one of the Cartesian axes. See if you can make the count 'count up' systematically; and 'count down' systematically. What happens to the count when the point (X,Y) is dancing back and forth across one axis? It's an up/down counter, and it should be alternating between two adjacent values on the display -- can you see why?

There are two kinds of adjustments you can make to the analog quadrature signals that may improve your up/down counting. One is to filter out high-frequency electronic noise from the signals; this is achieved by choosing a filtering time constant in the signal flow path. Your goal is *not* to filter out genuine optical signal variation, but only noise unrelated to fringe signals. To see the difference, go back to the breadboard-tapping mode, and see what happens to your locus if you pick too long a time constant.

The second electronic improvement you can apply is to choose a hysteresis level for the electronic discriminators that convert the analog signals to digital values. The goal here is to ensure that the digital switching occurs only on a genuine fringe signal, not on every little bit of noise and fluctuation as the analog signal passes through zero. Check Appendix K for a discussion of hysteresis, and how it's quantified, and how to see it in operation on your signal. Using either of the diagnostics mentioned in that Appendix, your goal is to make the hysteresis width rather greater than noise and fluctuations, but well under the amplitude of the genuine analog variation of the input to the discriminator. It may be that you need only minimal hysteresis for reliable counting in the quadrature mode.

When you have achieved reliable counting, note how many counts you get per full cycle of motion around your locus. If the relative phase of your quadrature signals really is 90°, then these counts come equally spaced in optical phase. To what size of path-length change does one single count correspond? What mechanical displacement of the mirror does this reveal? Yet smaller still than the least-count information on the counter is the information on your 'scope's display of the point on the locus: to what degree is the point's location stable? You might devise a test with your gas cell for changing the optical phase -- cap off the far end of a supply hose to the gas cell, and change the air pressure in the hose and cell merely by squeezing the hose. Every time you unpinch the hose, you should return to the same pressure; does the (X,Y) point return to the same place? To what degree? With what noise or drift level? If you can estimate the location around the periphery of the locus to one part in ten (or one 'hour on the clock', 1 part in 12), to what level are you estimating displacements in optical phase? What (one-way) mechanical motion of the end mirror would give the equivalent detectable signal?

To see another advantage of quadrature counting, let your interferometer settle down, and reset the counter to zero. If your draft cover is in place, and your temperatures have had time to stabilize, you can get a feeling for how fast (or slow) is their drift even by a quarter-locus or least-count. Now tap the breadboard once, watching the X-Y display to see what sort of tap it takes fully to 'fill in' the circumference of the locus during the vibration induced by your tap. Once the (X,Y) point on the locus settles back down after the tap, to what degree does it settle back down to the same place? And what has happened to the count on the counter? You might be able to see instantaneous counts, during the interval of vibration, that depart from zero-count by dozens or hundreds; and yet you might see that the count returns right back to zero, even after those long-lasting excursions a long way on either side of zero. If you can get this to work, find your favorite soft bouncy-ball and *bounce* it on your optical breadboard, and impress yourself, and any nearby opticians, with the amazing vibration immunity that quadrature counting affords you.

There are of course more quantitative uses for quadrature-counting capability. The combination of XY-display on a 'scope, and accumulated counts on the counter, give you the tools of choice for tasks like wavelength measurement, index-of-refraction measurement, and investigation of thermal expansion, the piezo-electric effect, and magnetostriction. These experimental capabilities are discussed in Sections 9-16, but you have completed the assembly of the optical instrument that makes all of them accessible and convenient to investigate.

4 Interlude on *Polarization*

All your work in interferometry to this point has made use of the wave nature of light, but has not really depended on whether light is a transverse or a longitudinal disturbance. But you know that light is, in fact, a transverse electromagnetic disturbance, and hence there remains the optician's choice over one additional degree of freedom, the direction of (say) its electric-field vector. This interlude will introduce you to a set of polarization techniques, valuable in general but also applicable to some really clever techniques in interferometry.

Your tools for this interlude include the HeNe and diode laser sources, the alignment towers as view screens, and (as diagnostic devices) the two Polaroid-type linear polarizers in rotary mounts on rod-and-post holder bases. In addition, you will get acquainted with the remarkably useful polarizing beam-splitter cubes (PBSCs) that you'll be using in other interferometers.

You might start with either of your lasers mounted in its support cradle, and turned on so as to deliver a beam across an otherwise empty optical breadboard. Send its beam to illuminate an alignment tower as view screen, and now interpose in the beam one of your Polaroids. Some light will still come through; show that you can *extinguish* the light by rotating the Polaroid. This is enough to show that your laser is delivering linearly polarized light. [Not all HeNe lasers do so -- generic HeNe deliver 'randomly polarized' light -- but yours does. Check your diode laser, too.] Now back to the extinction criterion -- what happens as you rotate the Polaroid away from that extinction angle? [here, meaning the rotating the Polaroid *in its own plane*, using the ring-mount that makes it easy to do so.] You should see more light, and you should see maximal transmission at a *pass* angle, ±90° away from the extinction angle.

To see that the polarization of the beam is produced at the source, now loosen the mounting screws in the laser holder, and rotate your HeNe laser bodily, by some modest angle, about its own light-output axis. Clamp it back down, and confirm that the extinction and pass angles of the Polaroid have rotated by the same amount.

Now set your lasers to a nominal condition of this rotation: with the HeNe, rotate until the power cord at the back end of the laser lies at the lowest position (6 o'clock on the dial) possible; with the diode laser, rotate so that the ceramic-faced thermoelectric modules contact the aluminum support below the laser. I claim that each beam now has the electric-field vector lying vertically in space -- but you need to confirm that. The method of choice is NOT to rely on the angular numbers printed on the Polaroid's mount; you can't know for sure if the optical pass direction of the Polaroid material has been aligned mechanically along the mounting ring, nor can you be sure whether it's the 0° vs. the 90° marking on its scale that's intended to mark the pass condition. Instead, it's time to perform the unambiguous 'Brewster test' to confirm the direction of your laser-beam's E-field.

To make this test, you need no Polaroids, but only a plain and uncoated slab of dielectric that you can handle; a microscope slide is conventional, but any small slab of window glass or even clear plastic will work fine. What you want to try is shown in the two diagrams below -- you want the beam to pass through the slab, and you want to identify the beams that originate from partial reflection of the beam from the faces of the slab. In Fig. 4-1 below, <u>a view from vertically above</u>, the reflected beam is being deflected by about 60° from its original direction, and is falling on a alignment tower for you to view. The diagnostic is to rotate the slab, about a vertical axis, so that the reflected beam swings through an arc of ±20° or so from this condition. What you're looking to see is if there is an orientation of the slab that causes the reflected beam's intensity to drop to nearly zero.

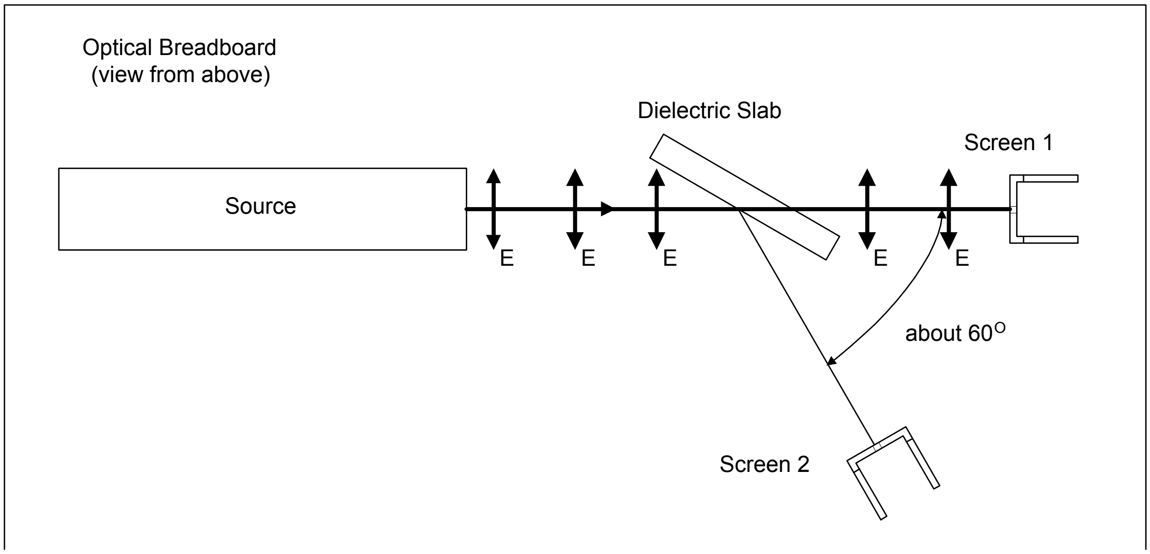

Figure 4-1: One way to try the 'Brewster test' for polarization – a view from **above**

In Fig. 4-2 below, you have <u>a side-on view</u> of another use of the same slab. Here the beam is being deflected downwards, through about 60°, and falling onto a white card that's been laid flat onto the optical breadboard. Now the diagnostic is to rotate the slab, about a horizontal axis, so that the reflected beam swings through an arc of ±20° or so from this condition. Again you're looking to see if there's an orientation of the slab that causes the reflected beam's intensity to drop to nearly zero. In *this* orientation of the slab, you should be able to see the deep minimum in spot intensity that develops when you've achieved the Brewster condition at the slab. If you can achieve a *zero* in reflected-beam intensity, then you have the slab oriented at the Brewster angle, *and* the laser beam with a vertical polarization.

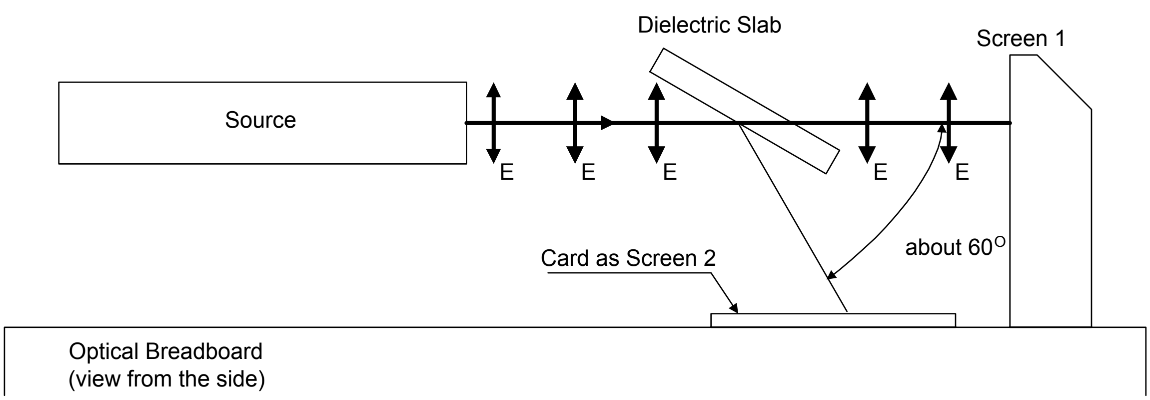

Figure 4-2: Another way to try the 'Brewster test' for polarization – a view from the **side**

[Rotate the HeNe laser source about its output-beam axis by perhaps 5-10° away from the previous orientation, and repeat the test in Fig. 4-2; you'll see a Brewster minimum again, but this time *not* getting right down to zero intensity. Exploit your understanding to get the laser oriented so as to deliver its E-field exactly vertically.]

Once you've accomplished this test, you can deliver a beam of confirmed vertical polarization (meaning, in modern optical usage, having a vertically-directed E-field), and now you can use it to test the meaning of the angular scale on your Polaroids. Be sure you write down some established fact, such as "For a vertically polarized beam, this Polaroid, set so its angular scale reads 1° or 181°, will pass the beam; but when set to 91° or 271°, will block the beam." Hereafter, your Polaroid will be the easy diagnostic for testing the polarization of optical beams.

Now rotate your HeNe laser by a full 45° about its long axis and confirm (by the extinction test) that it's delivering a linearly polarized beam (of tilted polarization, of course). Downstream, put a first Polaroid into position, set for now to its 0° setting. Now use a second Polaroid, downstream from the first, to find out

- whether the beam emerging from the first Polaroid is linearly polarized,
- and if so, what is the *direction* of its polarization.

Repeat these two tests with the upstream Polaroid set to its 90° setting.

You'll need to form a mental picture of what a Polaroid does: it does *not* merely change the intensity of a light beam passing through it, but *also* changes the light's polarization. In particular, by fully absorbing the light component with E-field perpendicular to its pass direction, it delivers an output beam whose E-field is wholly along its pass direction.

Now with your HeNe laser's E-field axis tilted to the 45° position, and a single downstream Polaroid available to rotate through its setting range of 0 - 90°, you can get a beam (admittedly of somewhat variable intensity) which is always linearly polarized, but whose polarization direction you can swing smoothly through the whole 0-90° range. Now let that beam fall onto one of your polarizing beam-splitter cubes (PBSCs) held in its upright holder and mounted on one of your beamsplitter bases. The optical element you see looks like a solid cube of glass, 1" on an edge, but it is in fact made of two 45-45-90° prisms glued together on a common hypotenuse. [Carefully remove the four mounting screws in the top plastic retaining plate to get a view of this common hypotenuse. While you're at it, note the two shallow grooves, one above and one below the cube, which hold the cube in place. Note also that the cube's hypotenuse is intended to lie parallel to the long axis of the rectangular top retaining plate, so you'll always know which orientation it has.]

A PBSC is intended for use with light falling nearly perpendicularly on its faces, and for such light, the internal hypotenuse face acts like a beam-splitter set at 45°. The result is that you might expect some of the incident light to emerge undeflected from the face opposite the entry face, and some of the light to emerge from a side face, partially reflected off the internal hypotenuse face. Set your PBSC in its mount on your breadboard so that your laser beam passes through it, and set up alignment towers to serve as viewing screens for these two beams. [Why don't you get any light emerging from the *fourth* face of the cube?].

Now look at the intensity of these two beams, as a function of the angular orientation of your upstream Polaroid. What does the cube do to an input beam that is vertically polarized? What does the cube do to an input beam that is horizontally polarized? What does it do to a generic beam?

You can do more than this. With the input beam at a generic angle of linear polarization, so that both output beams are in view, use your second Polaroid to diagnose the polarization of both output beams, the straight-through 'passed' beam and the deflected-by-90° 'bent' beam. You should find the 'passed' beam always has one state of polarization, whatever its intensity and whatever the state of the input polarization. You should also find the 'bent' beam always has another state of polarization, whatever its intensity and whatever the state of the input polarization. You should also know that the common-hypotenuse face internal to the cube is coated with a multilayer dielectric that absorbs a negligible amount of light, so that (*un*like a Polaroid) this PBSC device only redirects light, according to its polarization. You may think of it as 'sorting' light, which enters going one direction but with two polarization components, into two separate output beams, each of one single polarization (orthogonal to the other's).

If you use a z-axis to designate the vertical direction in your lab, and take as x-axis the direction of propagation of your input laser beam, then in this picture you can describe the general state of input light as

$$\mathbf{E}_{\text{in}} = (E_1\, \hat{\mathbf{y}} + E_2\, \hat{\mathbf{z}}) \exp i(k\, x - \omega\, t)$$

while the two output beams (at least for an ideal PBSC) can be written as

$$\mathbf{E}_{\text{pass}} = (E_1\, \hat{\mathbf{y}}) \exp i(k\, x - \omega\, t)\ ,\ \ \mathbf{E}_{\text{bent}} = (E_2\, \hat{\mathbf{z}}) \exp i(k\, y - \omega\, t)$$

Note that the 'bent' beam is propagating in a new direction, and note that each 'piece' of the input beam has been preserved in coefficient. In fact, with E_1 and E_2 standing for generic complex coefficients, the input beam need not even be restricted to a state of linear polarization. By virtue of the superposition principle for light, an input beam of *any* polarization has been disassembled into two beams, each preserving the amplitude of the original beam's components along **y** and **z**.

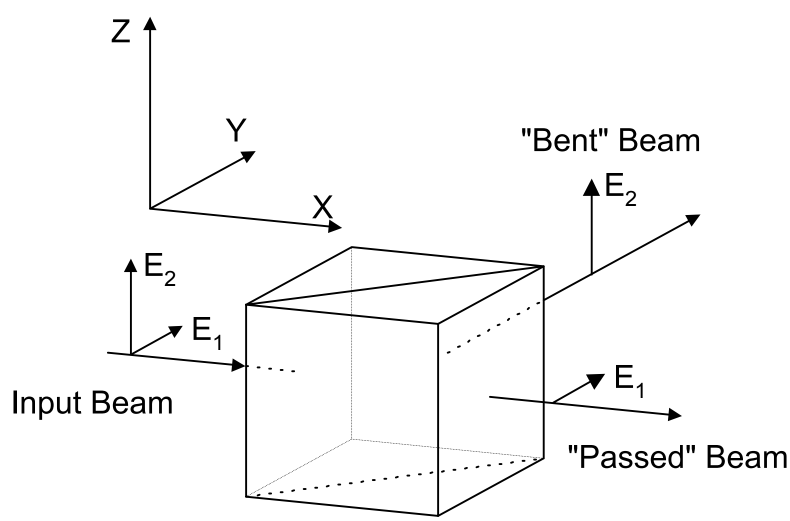

Figure 4-3: One choice of axes for understanding the polarizing beam-splitting cube (PBSC)

5 The Sagnac Interferometer

The first thing you need to know about a Sagnac interferometer is that the name is French, and the pronunciation is approximately 'Sahn-yock'. You'll build one, and find that it is not very well suited to measuring lengths, or length changes. You will find that it is an interferometer of unparalleled fringe stability, and you'll learn why. You will also see that it can be built in such a way as to exploit polarization of light in some truly clever ways, both in the two separated beams and in the detection of the light. Finally, you'll see that it's an interferometer of optimum 'photon efficiency'; in principle, every photon leaving the laser ends up on one of two detectors, and the output depends directly on the difference of these two detector signals. As a result, you'll get to experience an interferometer in which fringe stability well under *10^{-3} of a fringe* is readily attained.

a. The Sagnac topology

How is this all achieved? The first and characteristic difference in a Sagnac interferometer is that both of the two beams emerging from the beamsplitter end up going around all four sides of a rectangle (in opposite directions simultaneously), each returning to the original beamsplitter to re-combine there. The result is that any mechanical motion of any of the optical components is common to both beams, and (to first order) does not show up in the path-length difference at all. You'll find that it's actually possible to have the two beams traverse (in opposite directions) the very same path through the air, and contact the same points on the mirrors, further improving this 'common-mode rejection' of influences like mirror vibration and air turbulence.

The layout of a Sagnac interferometer that you can build with your TeachSpin components is shown in Fig. 5-1 below. There are a few special features of the components that are now worth describing.

a) The easiest source on which to learn the ideas is the HeNe laser, chosen for its good beam quality; because of the use made of polarization tricks, the HeNe is best rotated in its mount so as to deliver its linearly polarized output at an angle of 45° away from the vertical.

b) The beam reaches the interferometer proper via the two steering mirrors, making a near-90° bend at each mirror as usual. The second steering mirror is the one mounted on the one-dimensional rack-and-pinion slide, and that slide is oriented so that the beam leaving the second steering mirror can readily be translated sideways, remaining parallel to its original direction.

c) The interferometer starts with a polarizing beam-splitter cube (PBSC) serving as beamsplitter, and (given the state of input polarization) the beams emerging from it are approximately equal in power, orthogonal in direction, and also orthogonal in polarization.

d) Note that there are three mirrors that serve to define the other corners of a rectangle; these are mirrors of the sort used as end mirrors in the Michelson interferometer, but here are used at 45° angle of incidence. See Appendix L for the mounting holes to use in the mirror bases for this new angle of incidence.

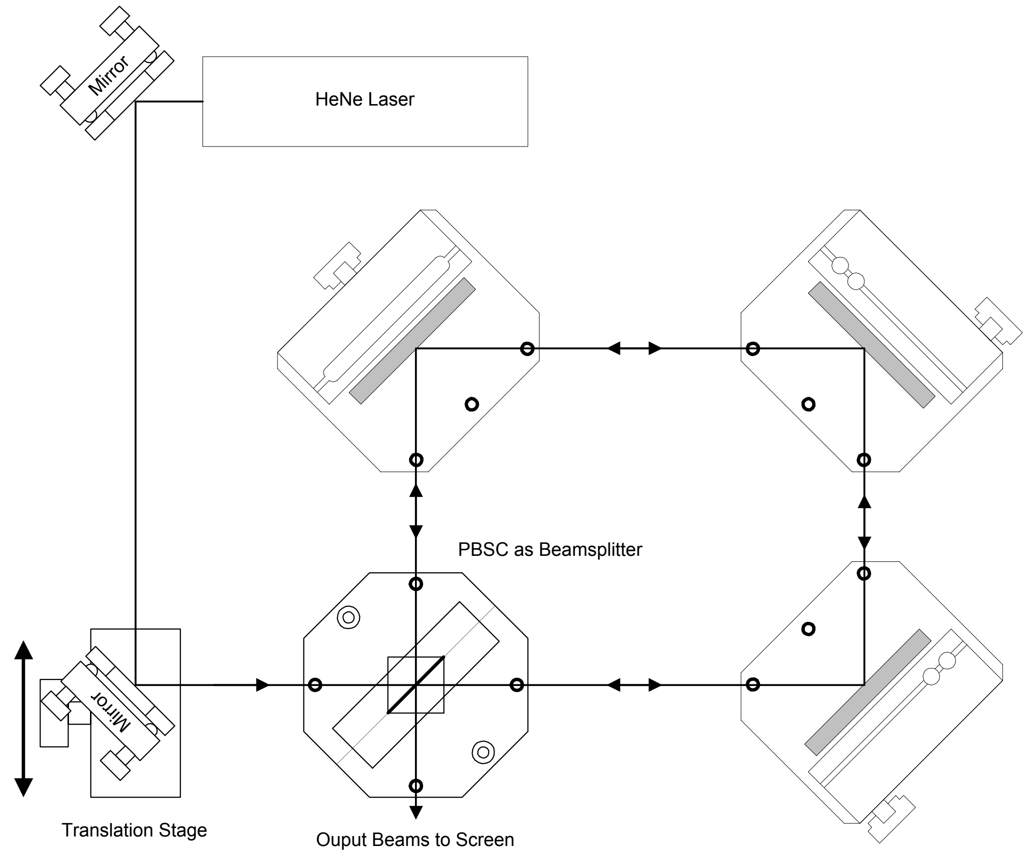

Figure 5-1: The topology of a Sagnac interferometer (not to scale)

e) After finishing a trip around the whole of the rectangle, each beam encounters the PBSC again, re-entering the cube at a different face from where it left the cube. Precisely because of the polarization-separating property of the PBSC, the input beam component which is ‘bent’ by the PBSC at its first encounter is also ‘bent’ by the PBSC at its return; and the input beam component which is ‘passed’ by the PBSC at its first encounter is also ‘passed’ by the PBSC at its return. The result is that *both* beams emerge from their second encounter with the PBSC at the heretofore-unused fourth face, and they head away from the PBSC in a conveniently accessible way.

f) To align the interferometer is to get the two counter-propagating beams to overlay each other in space all the way around the rectangular path, and also to overlay each other at the output. You’ll soon see how this alignment can be achieved.

g) The two beams emerging from the fourth face of the PBSC are still perpendicularly polarized, so they will *fail* to show interference fringes; that’s the motivation for a clever use of polarization optics in diagnosing the output of the interferometer.

b. Aligning a Sagnac interferometer

Now that you know the topology and the components required, here's a suggested procedure for setting up and aligning the interferometer. You'll need the two beam paddles, and you'll be mounting them at a variety of the convenient alignment holes that are found on the top face of the bases of the PBSC-holder and the mirror holders -- see Fig. 5-2 for reference to locations.

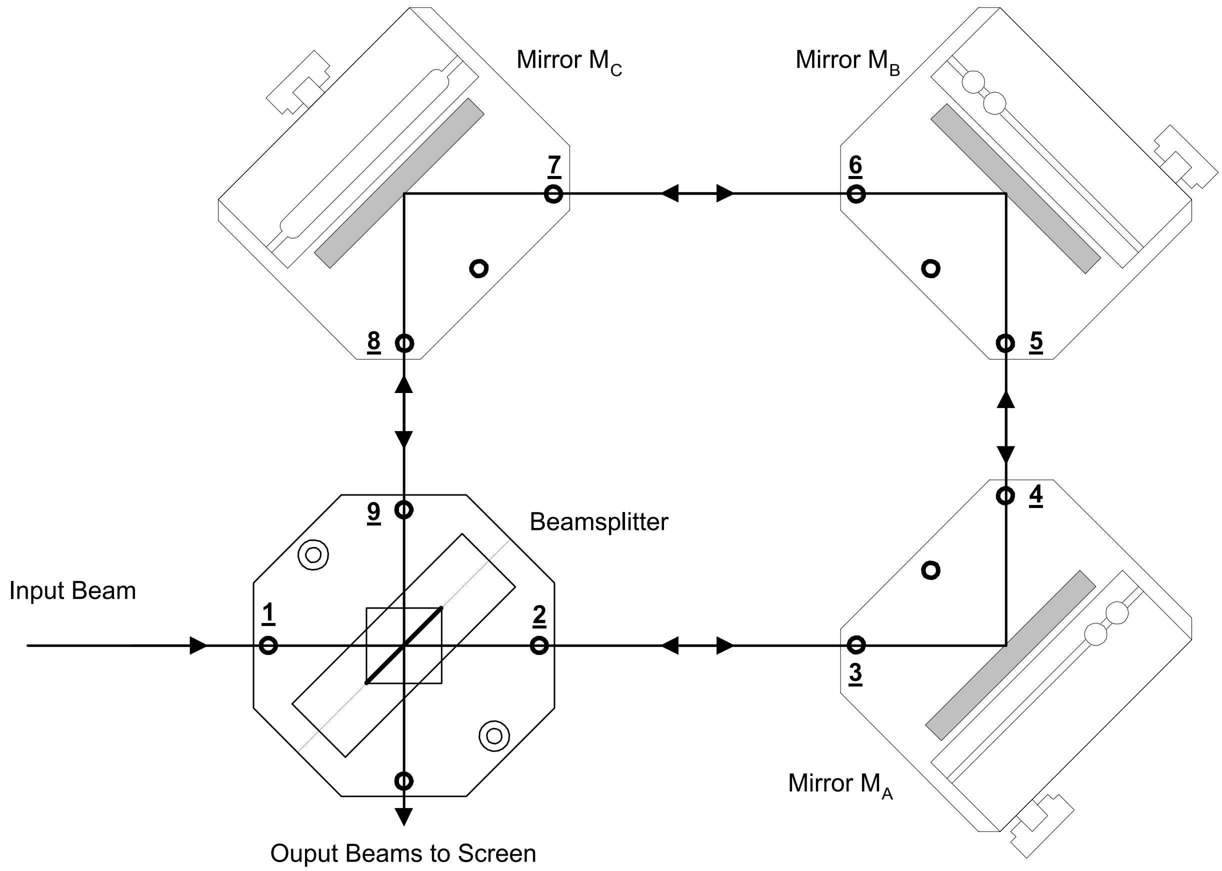

Figure 5-2: Labeling alignment-paddle locations in the Sagnac interferometer

1) Start with a HeNe laser and the two steering mirrors as introduced above, and then put down onto your optical breadboard the PBSC on its base, and mirror M_A as shown. Note that M_A gets the 'passed' beam from the PBSC, and note that M_A is mounted (now using different screw holes in its base) so as to reflect the laser beam through a 90° angle. For now, it would be well to block the 'bent' beam from the PBSC with an alignment tower, since your first goal is to use the steering mirrors M_1 and M_2 to get the 'passed' beam to pass through beam-paddles placed at locations #2 and #3 as shown in the Figure.

2) Now remove the paddles, block the 'passed' beam with an alignment tower, and let the 'bent' beam reach mirror M_C, also mounted so as to deflect a beam through 90°. [It is important that mirrors M_A and M_C have orthogonal directions of the 'hinge lines' in their flexure hinges.] Put beam paddles in place at locations #8 and #9, since you'll want the 'bent' beam to pass through holes in them. But don't use the steering mirrors M_1 and M_2 to

achieve this; instead, achieve the aiming you require by loosening the bolts holding the PBSC-holder to the optical breadboard, and using the ‘slop’ in its mounting holes. A rotation about a vertical axis of the whole PBSC base will give you one degree of freedom for centering the beam on the beam paddle at location #8; the other perpendicular degree of freedom needed can be had from shimming under one corner of the PBSC mount’s baseplate. [See Appendix G for what ‘shimming’ means in this context.]

3) Now you have beams reaching mirrors M_A and M_C, and you have those mirrors in the right positions and approximately at the right angles. You should now put mirror M_B in position to complete the rectangular perimeter of the interferometer, and you should see two beams hitting its front face. Bolt it loosely to the breadboard, also at the 45° orientation, and mount beam paddles on its base at locations #5 and #6. Now loosen the mounting of M_A and rotate its base on the breadboard until it sends a beam through the paddle at location #5; similarly loosen the mounting of M_C and rotate its base on the breadboard until it sends a beam through the paddle at location #6. Now if you remove the two paddles, you should see one *single* spot on mirror M_B that is illuminated by two separate beams from the PBSC. You will be able to use a fine-adjustment screw on the back of either M_A or M_C to help with the vertical degree of freedom in this overlap. Once you’ve achieved this, you can tighten down mirrors M_A and M_C.

4) Before tightening down mirror M_B, use a beam paddle inserted at location #4 or #7 to identify *beams reaching it from both sides*. You might temporarily use a sheet of paper instead, since with such a semi-transparent beam indicator you get a good idea of the relative position of two beams reaching the paper from opposite sides. Your goal is to rotate mirror M_B by the ‘slop’ available in its breadboard-mounting holes until you’ve achieved overlap of the two beams hitting the opposite sides of the paper sheet; you may also use the thumbscrew adjustment on the back of mirror M_B. Don’t forget at this stage to tighten down the mounting screws of M_B, and indeed any others that may still be loose.

5) Now put an alignment tower in position to see the *two* beams that emerge from the fourth face of the PBSC. They will probably not yet be overlapping, but they should both be emerging from near the center of the face of the PBSC. If you now adjust the thumbscrew on the back of M_A or M_C, you will find that turning *either* thumbscrew affects the position of *both* spots on your alignment tower. No matter; you can still use those two thumbscrews to give a good overlap. [You should not need to use the thumbscrew on M_B to achieve this overlap.]

6) But when you’ve achieved overlap, you *won’t* see any fringes on your screen. To see why not, readjust the spots so they’re a bit away from overlapping, and interpose a Polaroid between the PBSC and your view screen to diagnose the polarization of the two output beams. You should find that the two beams are perpendicularly polarized, and you should mentally trace each beam backwards until you can understand where it got its polarization. You should also understand the theory of why two orthogonally polarized light beams do not show interference phenomena. Finally, you should orient your Polaroid at 45°, and see that both beams now are transmitted (in part) by this Polaroid.

7) Now repeat the adjustments that bring the two spots into overlap, and you should see fringes. They'll be *great* fringes, of superb contrast and stability. The contrast is high because the two 'arms' in this interferometer have nearly identical lengths, so you're optimally set for mutual coherence; they're stable for the reasons described under 'common mode' above, and the fringes' geometrical stability keeps the contrast high in your view on the screen. The final stage of adjustment is not to perfect the overlap of the beam spots, but instead to lower (to zero) the 'spatial frequency' of the fringes. See Appendix M for a discussion of what gives rise to these straight-line fringes and their periodicity on the screen.

8) The very stability of the Sagnac interferometer makes it hard to know where on the interferometer you can 'push' to move through a succession of fringes. So now you can exploit a wonderful feature of this interferometer, by moving *away* from the present configuration that has the two beams nearly perfectly overlapping in space all the way around the interferometer. Go ahead and put beam paddles back in at locations #2 and #3, and now go back to the steering mirror M_2 and (finally) exercise its rack-and-pinion slide capability. You'll want to translate the beam now traversing the *central* holes in paddles at #2 and #3 until it passes instead through the *side* holes in the same paddles. [It doesn't matter to which side you go, but do make it the *same* side on both paddles.] You may need to trim up the adjustments on steering mirrors M_1 and M_2 to align the beam with the paddles; it might help to block the beams between locations #8 and #9 to isolate the beam you are using for alignment.

9) Once you have the input beam thus translated, you should remove all the paddles, and you should immediately see the two output beams, passing through the Polaroid, overlapping on the screen. Again you can touch up the thumbscrews on M_A and M_C for optimal fringes, and again you'll have the Sagnac interferometer working. The difference is that now at any point around the interferometer, the two counter-propagating beams are horizontally separated in space, like the two lanes of traffic on a two-way street. Confirm this by using a sheet of paper as a viewing screen; confirm that you understand the direction of propagation of both of the beams you see; also diagnose the *polarization* of the two beams you see.

10) The value of having the two beams separated in space is that you can now do something to one beam, separate from what you do to the other beam. The TeachSpin gas cell is designed so that it can be mounted with one of the two separated beams passing through it, and the other passing beside it. So now use the hose and syringe to make small pressure changes to the air inside the cell, and see what that does to the intensity of the spot on your view screen. You should see a sinusoidal variation in brightness.

11) Now that you're persuaded that you can make fringes go by, it's time to use a more clever and efficient detection scheme than the Polaroid-and-view screen combination. You want to remove the Polaroid, and to use instead the second PBSC in the curiously tilted arrangement shown in Fig. 5-3; here the overlapped output beam from the Sagnac encounters a PBSC oriented at 45° relative to the horizontal and vertical. [See Appendix N for the procedure required for moving a PBSC from one of your upright beamsplitter mounts into this tilted-cube mount.] For starters, it's well to leave the photodetectors out of the holders that will soon accommodate them; instead, use a large white card that will give you a view of

the beam ‘passed’ by the second PBSC. Clearly, this beam, like the beam that passed through the Polaroid, represents the overlap and interference of the two beams going around the interferometer, so you should see fringes at this spot. Now note that the beam ‘bent’ by the second PBSC is further deflected through 90° by an ordinary mirror, so that it too comes out to your viewing card, propagating parallel to your ‘passed’ beam but displaced ‘up and over’.

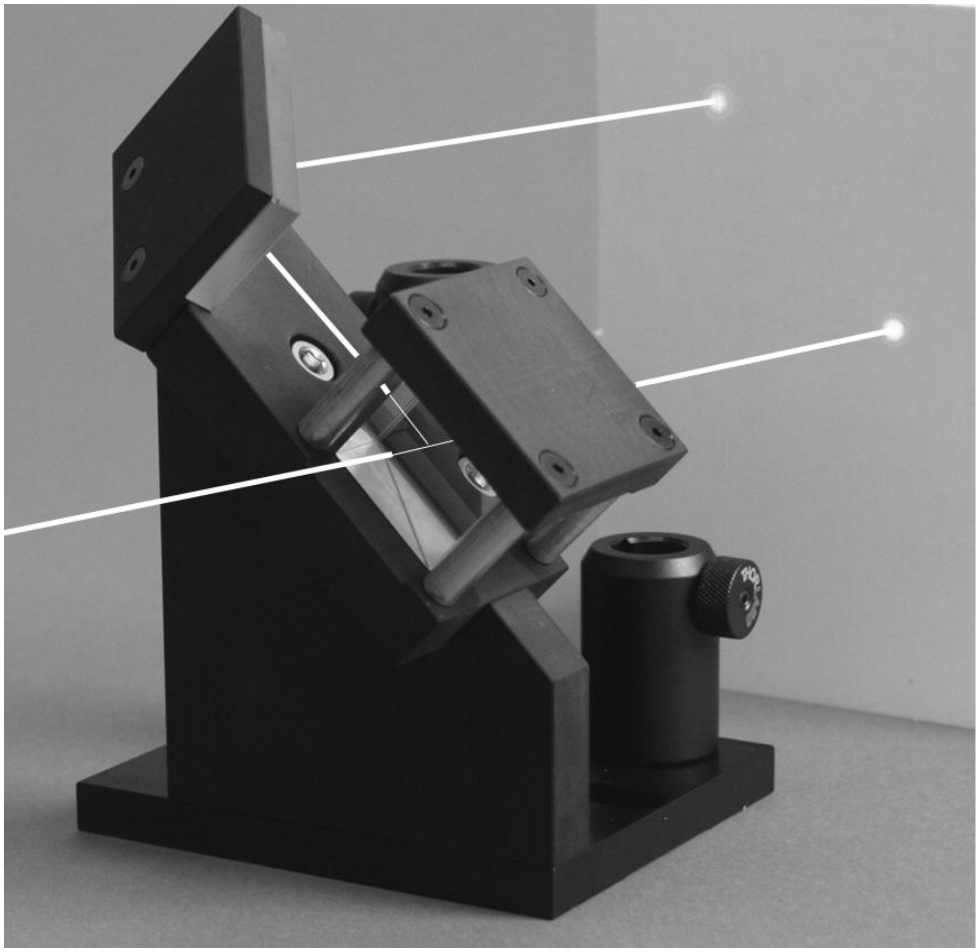

Figure 5-3: Using the tilted PSBC as a polarization separator

12) Your goal is now to adjust your interferometer for zero spatial frequency at both spots, and then to notice that the light intensities in the two spots on your card are complementary -- that one is brightest when the other is dimmest. The reason for this is laid out mathematically below, but it is also obvious on physical grounds: all the light from the laser enters the first PBSC; none of it is absorbed there, so all of it is sent around the interferometer. That light, reaching the first PBSC again as beam recombiner, emerges from the ‘fourth face’ of the first PBSC, and heads to the second PBSC. That element is also lossless, so all the light from the laser has to emerge into one of the two spots on your card. [Naturally there are losses in the interferometer, due chiefly to reflectivity less than 100% at the metal front-surface mirrors in the interferometer; but in principle this interferometer and

this polarimetric detection scheme can be made to be 100%-efficient: all the light leaving the laser would end up on one detector or the other.]

13) Now it's time to put photodetectors in place to quantify these two output beams. Use the rod-and-post holder arrangements to get the two detectors at the right heights; slide the whole polarimetric-detector baseplate laterally to get the horizontal adjustment you need for the 'passed' beam; and use the lateral slide adjustment of the taller post's base to get the horizontal adjustment you need for the 'bent' beam. When you have both beams hitting near the center of their detectors, power up both detectors, and set them to equal gains that also keep their readings on scale. Set up a dual-trace oscilloscope to view the two intensity reading simultaneously, and sweep at perhaps 1 s/division to see what happens to the two signals as you vary the optical phase difference (via the gas pressure in the cell). You should see the complementarity of the two signals; you should also note that the amplitudes of the two sinusoids are very nearly equal.

14) Given two such analog electrical signals, it's a very clever technique to subtract one from the other. That will give a signal which varies on both sides of zero, rather than lying always on one side of zero. The difference signal will then have 'zero crossings', arising when the two beam spots on a card would be of equal intensity. One of the many attractions of such a 'zero crossing' signal is that its location is first-order independent of any power fluctuations in the laser. You can play with the 'subtract' function you might have on your 'scope, but you can also use the signal-flow schematic on the front face of the Modern Interferometry controller to perform this analog subtraction. There is provision to match the amplitude of the two sinusoids, in case detector sensitivities are not quite matched; there is also provision to filter out high-frequency noise in the signals, leaving the steady optical-phase-difference signal controlled by the gas cell. You should begin to see that the interferometric difference signal has *amazingly* low noise, even after you use the further amplification of the difference signal that your controller box permits.

15) Since the two beams in the interferometer are separated in space, there is some non-zero sensitivity to air-density fluctuations inside the interferometer. So this might be the time to re-install the draft shield over the rectangular perimeter of the interferometer. Now you can bring the diagnostic difference signal down to a zero-crossing, park it there, raise the gain on your 'scope detection of the result, and begin to see just how stable a signal you can get from an interferometer. Go ahead and tap the optical breadboard to get an idea of how well vibration is an effect common to both 'arms' of this interferometer, and hence cancelled out. See if you can tell what limits the stability in time of your signal; you may find that time-varying pressure in the gas cell is the biggest problem.

16) There are many ways (other than the use of the gas cell) to vary the optical phase difference in the two arms of the interferometer. One of them is based on the index of refraction of thin slabs; find the holder which carries *two* glass slabs of 1-mm thickness and 1-cm^2 area, and find the angular rotation stage which can rotate its mounting post about a vertical axis. The glass slabs are laid out so that each is canted, or tilted, by 10.0° (±0.5°) away from face-on in their holder. They are also arranged such that you should be able to

position the whole unit (glass slabs, post and rotator) so as to have the two separated beams in the interferometer pass through the two slabs of glass.

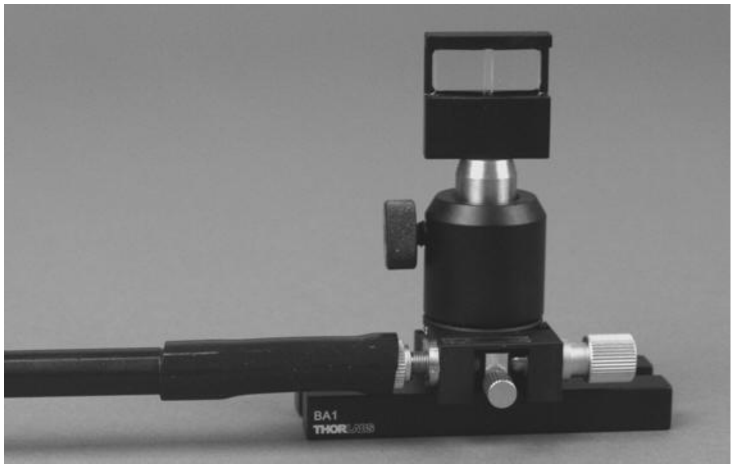

Figure 5-4: The rotation stage, bearing a two-plate tuner for the Sagnac interferometer

When you've gotten this unit into position, clamp it down to the breadboard, and then use thumbscrews at M_A and M_C to trim up the fringe patterns in the two output beams. Now you'll find that as you rotate the glass-slabs' holder, you get a steady succession of fringes, but one which you can 'park' at any desired point with superb stability. The full description of why this occurs is left for later, but note that if the two slabs start with each one illuminated by light making a 10° angle of incidence, then turning the rotator by (say) 1° changes the angle of incidence for one slab to 11°, and the other slab to 9°. Clearly this increases the effective thickness of the slab in one arm, and decreases it in the other, giving rise to the phase difference you can detect.

17) The separated-beam Sagnac interferometer offers access to the two individual beams in it, but it thereby becomes somewhat sensitive again to two noise-contributing effects. For one, the air density will inevitably be somewhat different at the two distinct beam locations. For another, the two beams are now interacting with different places on the three bending mirrors, so that vibrations of the mirrors can affect one beam differently than the other. So you may want to remove the 'two-plate glass tuner' and use the rack-and-pinion adjuster at steering mirror M_2 to take you back to the fully overlapped-beam condition; you might temporarily put paddles back in at locations #1 and #2 to align the beam 'passed' by the PBSC with your chosen axis, and then trim up the fringe pattern by adjusting corner mirrors M_A and M_C. [Having translated the input beam sideways, you'll also need to translate bodily sideways the whole polarimetric detection unit, by its base, to get it centered on the new location of the output beam.] You're back to having the two beams in your interferometer (nominally) overlapping each other all the way around the perimeter of the interferometer, which helps to defeat both noise sources just mentioned. But if also defeats the use of the two-plate tilter to tune the phase difference between the two beams. You might think of trying the one-plate glass-slab tilter instead, but that plate will cause a phase shift that should be common to both beams, so it still won't work as a phase adjuster.

So here's the trick: you also have a one-plate tuner, but this one made of a slab (0.5 mm thick) of crystalline quartz. It is a plate with two distinct indices of refraction, which we may label n_h (for horizontally polarized light) and n_v (for vertically polarized light). The value of this plate is that the two beams traversing your interferometer are not only going in opposite directions; they are also of perpendicular in polarization, and so the two beams will sample the two distinct indices of refraction and thus encounter effectively different plates. So set up the quartz one-plate tilter unit somewhere inside the interferometer, and confirm that rotating it will shift the electronic difference signal you've been detecting. The goal is not to scan through a huge range of fringes, but rather to be able to bring the interferometer from any given starting condition to one of the zero-crossing points of the difference signal.

After confirming that it works, you can take the quartz tilter plate out of the interferometer, and try re-installing it between steering mirror M_2 and the beam-splitter PBSC, or between the PBSC and polarimetric detection unit, and confirm that it still does its job, and try to explain why.

Once you're able to put a draft cover over an overlapped-beams Sagnac interferometer, and once you're able to use the quartz tilter to reach a zero-crossing signal, you can use electronic gain as desired to look at the stability of your interferometric signal. How noisy is it in the short term? How stable is it in time? If there were to be a square-wave modulation superimposed on the signal you see, how large would it have to be to be distinctly detectable? How large a square-wave phase difference would produce this square-wave difference signal? You are now ready for Section 15, in which your newfound capabilities are used to detect the electro-optic effect in a solid material.

18) You might now check the degree to which you have the counter-propagating beams in your interferometer actually overlapping in space. Use a thin-paper viewing screen to get a view of beams impinging on both sides simultaneously, and test the beams near the PBSC (where the overlap should be near-perfect) and near bending mirror M_B (where the overlap might be at its worst). If the two beams are separated, or of visibly imperfect overlap, near M_B, then you have a diagnostic against which you can improve. The payoff will be lower sensitivity to air density and vibration-induced noise in your output signal; the method is to pay more attention to aligning out the imperfections of the various mechanical and optical components in your interferometer.

c. Understanding polarimetric detection

Here's a computation intended to show how the output of the Sagnac interferometer can be represented mathematically, and how the tilted PBSC detector unit works with that optical output to generate the electrical signals you have seen. We imagine that the input of the Sagnac interferometer is linearly polarized light with direction of polarization tilted at 45°, so that there are vertical and horizontal electric-field components of equal magnitude and in phase:

$$\mathbf{E}_{in} = \frac{E_0}{\sqrt{2}}\,\hat{\mathbf{y}}\exp(ikx - i\omega t) + \frac{E_0}{\sqrt{2}}\,\hat{\mathbf{z}}\exp(ikx - i\omega t) \ .$$

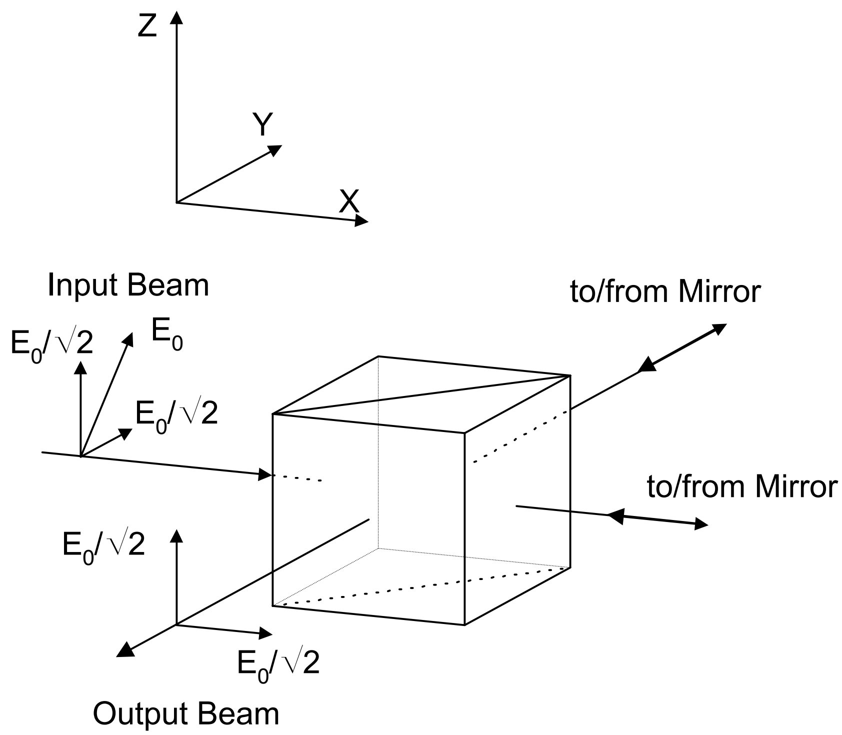

Figure 5-5: Input and output states of polarization for a PSBC

Now the beam emerging from the beamsplitter PBSC is propagating in a different direction; we'll assume that the beam's passage through the interferometer has not changed the amplitude of either component, but has given two distinct phase shifts (to the clockwise- and counterclockwise propagating beams):

$$\mathbf{E}_{out} = \frac{E_0}{\sqrt{2}} \hat{\mathbf{y}} \exp(i\,\phi_{ccw}) \exp(-iky - i\omega t) + \frac{E_0}{\sqrt{2}} \hat{\mathbf{z}} \exp(i\,\phi_{cw}) \exp(-iky - i\omega t) \ .$$

This beam is in a complicated state of polarization (elliptical in general), but it's not too hard to model what will happen to it upon encountering the second, tilted, PBSC. The projection of $\mathbf{E}_{out}$ along the vector direction $\mathbf{n}_1$ will be the part of $\mathbf{E}_{out}$ that will emerge from one of the faces of the second PBSC, while the projection of $\mathbf{E}_{out}$ along the direction $\mathbf{n}_2$ will emerge from the other face of this PBSC:

$$\mathbf{E}_{\#1} = \mathbf{E}_{out} \cdot \hat{n}_1 = \mathbf{E}_{out} \cdot \frac{\hat{\mathbf{x}} + \hat{\mathbf{z}}}{\sqrt{2}} \ ; \mathbf{E}_{\#2} = \mathbf{E}_{out} \cdot \hat{n}_2 = \mathbf{E}_{out} \cdot \frac{-\hat{\mathbf{x}} + \hat{\mathbf{z}}}{\sqrt{2}} \ .$$

Since these describe the electric fields in two separate beams that are then carried to the two distinct detectors, we want a model for the power in the two beams. In the complex representation of an electric field, that is given by

$$P = \frac{1}{2}|\mathbf{E}|^2 \ ,$$

and the signal that we ultimately observe electrically is given by

$$P_1 = \frac{1}{2}|\mathbf{E}_{\#1}|^2 \ ; P_2 = \frac{1}{2}|\mathbf{E}_{\#2}|^2 \ ; \ P = P_1 - P_2 \ .$$

There is a fair amount of algebra required to compute the details of these signals, but the results are

$$P_1 = \frac{1}{2}\frac{E_0^2}{4}(2 + 2\cos\Delta\phi) \;\; ; \;\; P_2 = \frac{1}{2}\frac{E_0^2}{4}(2 - 2\cos\Delta\phi) \;\; ,$$

where $\Delta\phi \equiv \phi_{ccw} - \phi_{cw}$ gives the phase difference accumulated by the two beams making their way, in opposite directions, around the ring of the Sagnac interferometer. The significance of the leading coefficients is best seen by noting that

$$P_1 + P_2 = \frac{1}{2}\frac{E_0^2}{4}(4 + 0) = \frac{1}{2}E_0^2 = P_{in} \;\; ,$$

so sure enough, the two beams headed toward the detectors carry collectively all the power in the beam sent into the interferometer. To put it another way, we can see that the quantity P_{in} also represents the maximum power, alternately occurring in beam #1 and beam #2. Hence if we change from optical powers P to electrical signals S, we can predict that the signals on the two electrical outputs will behave according to

$$S_1 = \frac{S_{max}}{4}(2 + 2\cos\Delta\phi) \;\; ; \;\; S_2 = \frac{S_{max}}{4}(2 - 2\cos\Delta\phi) \;\; .$$

But in actual use, we view not the sum, but the difference, of the two signals, which gives a signal taking on both signs,

$$S(\Delta\phi) \equiv S_1 - S_2 = \frac{S_{max}}{4}(4\cos\Delta\phi) = S_{max}\cos\Delta\phi \;\; .$$

Here the quantity S_{max} has a very readily understood empirical meaning, since S_{max} gives the maxima, and $-S_{max}$ gives the minima, of the electrical signal you've directly monitored.

This also displays the sensitivity of the interferometer to changes in the phase difference $\Delta\phi \equiv \phi_{ccw} - \phi_{cw}$. We need only compute the derivative

$$\frac{\partial S}{\partial(\Delta\phi)} = S_{max} \cdot -\sin\Delta\phi \;\; ,$$

which in turn tells us that the output signal S will change by amount

$$\delta S = \pm S_{max}\,\delta(\Delta\phi)$$

if we choose to operate near the zero-crossing points of the signal. To be concrete, we suppose that S_{max} = 8 Volts, and that we change the phase difference by one one-*thousandth* of a full fringe, ie. $\Delta\phi = 0.001 \cdot 2\pi = 0.00628$ rad. Then this equation predicts a change in the output signal of size

$$\delta S = (8 \text{ Volts})\,(0.00628 \text{ rad}) = 0.050 \text{ Volts} = 50.\text{ mV}.$$

Whether this sort of 'one milli-fringe' change is detectable depends on the noise level of the interferometer's output signal; but some attention to vibration control, and the use of the draft shield and the fully-overlapped beam mode of the interferometer, ought to give you a noise level of only a few mV. Hence you might be in a position to detect a one milli-fringe signal with a signal-to-noise ratio better than ten!

6 Interlude on *Relativity*

Your textbook acquaintance with interferometry probably comes from the mention of the Michelson interferometer in connection with the Michelson-Morley experiment, which famously *failed* to detect the absolute motion of the earth with respect to the 'lumeniferous ether' in which light was then supposed to propagate. This introduces the tight connection, experimentally and theoretically, between relativity and the Michelson and Sagnac interferometers.

Here's the calculation that motivated Michelson and Morley. If there is a unique 'ether' frame in which light really propagates at speed c, then classical notions of space and time would suggest speeds of $c \pm v$ for the propagation of light with respect to an apparatus which was itself moving at speed v relative to the 'ether frame'. So a conventional distance of L, lying parallel to the direction of that velocity, would be traversed in a round trip time T_{para} most easily computed in the frame in which the mirrors are fixed, but through which the 'ether' is moving:

$$T_{para} = T_{out} + T_{back} = \frac{L}{c+v} + \frac{L}{c-v} = \frac{2L}{c}(1 - v^2/c^2)^{-1} \quad .$$

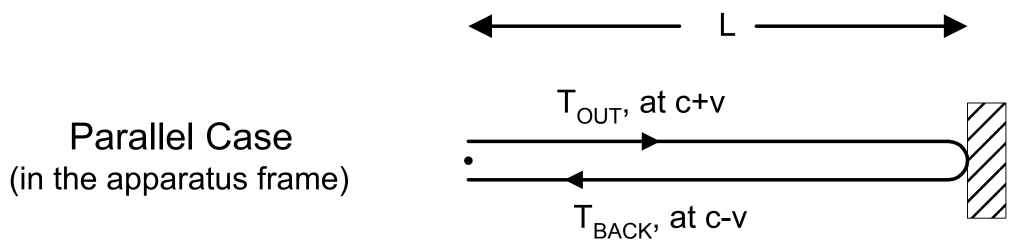

Figure 6-1: Distances and speeds for light paths (parallel to the motion) in a Michelson interferometer, under the theory of lumeniferous ether

Meanwhile, light in the perpendicular direction has a transit time best computed in the 'ether frame' in which propagation is always at speed c, but through which frame the mirrors are moving. In this frame, light would have to take on a path of triangular shape (relative to the ether) to return to its starting point in the apparatus, with a flight time of T_{perp} obeying

$$c\, T_{perp} = 2\sqrt{L^2 + (vT_{perp}/2)^2} \quad \text{, giving} \quad T_{perp} = \frac{2\,L}{c}(1 - v^2/c^2)^{-1/2} \quad .$$

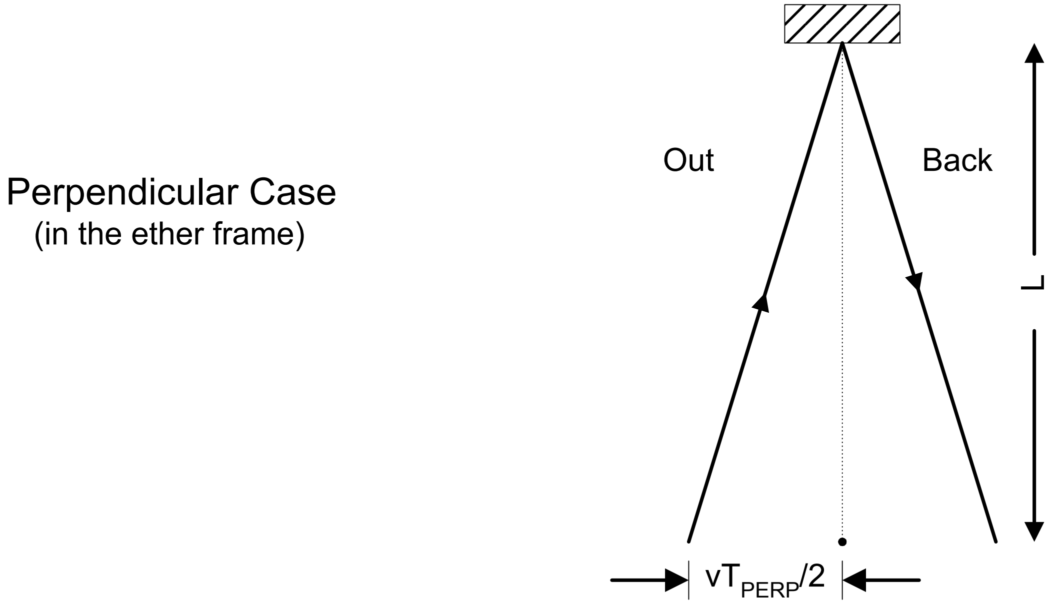

Figure 6-2: Distances and speeds for light paths (perpendicular to the motion) in a Michelson interferometer under the theory of lumeniferous ether

These two round-trip times *differ*, by an amount best seen by expanding in the small quantity v/c:

$$T_{para} - T_{perp} = \frac{2L}{c}\left(1 + 1\frac{v^2}{c^2} + \ldots\right) - \frac{2L}{c}\left(1 + \frac{1}{2}\frac{v^2}{c^2} + \ldots\right) \cong \frac{2L}{c}\frac{v^2}{2c^2} \quad .$$

[The relativistically astute reader will note that travel times computed in two distinct reference frames are being compared here; but of course before 1905 everyone was convinced of the existence of 'absolute time', and never dreamed of the importance of this point.]

It's conventional to regard the orbital motion of the earth around the sun, at speed v = 30 km/s, as the plausible speed of the apparatus relative to the ether frame; this gives $v/c = 1 \times 10^{-4}$ and $v^2/c^2 = 1 \times 10^{-8}$. One then wants to see this tiny fractional change in a time-of flight which is itself very small indeed: for an apparatus of scale L = 1 m, the one-way time-of-flight of light is $L/c = 3 \times 10^{-9}$ s. The time difference in the two times-of-flight is then

$$\Delta T = \frac{2L}{c}\frac{v^2}{2c^2} = \frac{L}{c}\left(\frac{v}{c}\right)^2 = (3 \times 10^{-9}\ \mathrm{s})\,(10^{-4})^2 = 3 \times 10^{-17}\ \mathrm{s} \quad ,$$

and appears to be inconceivably too small to detect. It is to Michelson's credit that he realized that interferometry potentially possessed the required measurement sensitivity; recall that a round-trip phase delay of 2π radians will give a whole fringe of signal, so a round-trip time delay of one period of visible-light's oscillation (just about 2×10^{-15} s) will give a full fringe of detectable signal. So the effect predicted in the ether theory is over one one-hundredth of a fringe, not beyond the realm of the detectable in a plausible apparatus. [You might look up a report on the original Michelson-Morley experiment to see what size of interferometer they used, and why they mounted it on a table they could rotate. You might also compute the actual fringe shift you'd predict, according to this line of thinking, for a Michelson interferometer you could actually build on your breadboard.]

Now the predicted fringe shift was definitely *not* seen, and there are lots of explanations for why it wasn't. An early explanation was in terms of the 'Lorentz-Fitzgerald contraction', which pointed out that on the basis of plausible theories of the composition of matter, the electric forces binding together the table would *themselves* change with its motion, and would cause a actual mechanical contraction (of only the *parallel* leg of the interferometer) by an amount

$$\Delta L = L\sqrt{1 - v^2/c^2} - L \cong L\left(-\frac{1}{2}\frac{v^2}{c^2} + \ldots \right) ,$$

that is just enough to eliminate the expected fringe shift. A deeper explanation was provided by the Special Theory of Relativity, which hypothesized the very nature of space and time was such that no experiment, mechanical or optical, would be able to detect uniform translational motion at any speed.

The Michelson-Morley experiment has since been repeated using other laser methods, and has reached a sensitivity vastly greater than that which would be needed to detect the effect of the earth's translational motion, if it indeed had effects of the size computed above. The null detection is so dramatic that some persons have gone on to suppose that absolute motion must of its very nature remain undetectable out of mere philosophical necessity. This reasoning has it that motion can only be defined with respect to something else (another body, or the ether, or space-time), so that absolute motion is a meaningless concept.

The fallacy in this plausible line of argumentation has been laid bare in a wonderful passage by Richard Feynman, who introduces his treatment of relativity in a section entitled "Relativity and the Philosophers". In particular, he points to the readily detectable existence of various effects due to uniform *rotational* motion, philosophical arguments to the contrary notwithstanding. Nothing can replace your reading the whole brief section [found in *The Feynman Lectures*, vol. 1, ch. 16, pp. 16-1 to 16-3]. But this does bring up the opportunity to discuss the Sagnac interferometer, whose original motivation was just this possibility of using optical means to detect 'absolute rotation'.

It's easy to do a naive calculation of the time difference expected for trips both ways around a closed path, and easiest of all to imagine counter-propagating circular paths or radius R in an instrument rotating about its center at angular velocity Ω relative to 'absolute space'. In a time T, an apparatus will rotate by angle ΩT, so the time T_{cw} it will take light leaving a source to return to the source-point again, as destination, is given by

$$T_{cw} = \frac{L}{c} = \frac{(2\pi + \Delta\theta)\,R}{c}, \text{ where } \Delta\theta = \Omega\, T_{cw} \quad ;$$

this can be solved to give

$$T_{cw} = \frac{2\pi\, R/c}{1 - \frac{\Omega R}{c}} \quad .$$

The time for counter-clockwise travel of the light around the same path, from starting point to reach the (rotated) source-point again, is given by

$$T_{ccw} = \frac{2\pi\, R/c}{1 + \frac{\Omega R}{c}} \quad .$$

The time difference can be extracted, and it is of *first* order in Ω:

$$\Delta T = T_{cw} - T_{ccw} = \frac{2\pi R}{c} \frac{2\,\Omega\,R}{c} = \frac{4\pi R^2}{c^2}\,\Omega = \frac{4\,A}{c^2}\,\Omega\ ,$$

where the answer is now given in terms of the area A enclosed by the counter-propagating paths, and where the result is in fact valid for any arbitrary common shape of the paths. The size of this time delay is not that small, and it renders the rotation of the *earth itself* detectable for a big enough interferometer. Via the Sagnac effect, optical gyroscopes are capable of detecting absolute rotation, even though Michelson interferometers are incapable of detecting absolute translation.

The justification of the naive derivation above in terms of special, or general, relativity is naturally much harder, and any discussion of the detectability of absolute rotation is sure to stray far from optics into 'centrifugal effects' and even to Mach's principle. But the mere fact of the existence of working optical gyroscopes is enough to underscore the importance of Feynman's point: in the face of the apparent *un*detectability of 'absolute translation', the detectability of 'absolute rotation' is "a question that can be answered only by experiment".

Meanwhile, you might compute the area A enclosed by the Sagnac interferometer you have built, and also compute the sort of time delay that you could potentially detect as a non-zero signal, using the relationship between detectable phase shift and inferred differential travel time:

$$\frac{\Delta\phi}{2\pi} = \frac{\Delta T}{T_{light}} = \frac{\Delta T}{\lambda/c}\ .$$

You should be able to show that quite modest rotation rates of an optical table are predicted to give time delays, and phase shifts, which you have the sensitivity to detect. [The question of *systematic* effects, and various experimental practicalities, we leave to ambitious experimenters.] In a tabletop experiment you'd have the rotation rate Ω under your control as an independent variable, and thus the ability to see how the phase shift depends on your choice of Ω. You'd also have the advantage of being able to separate Sagnac effects of order Ω from centrifugal effects of order Ω^2. The detection of the *earth's* rotation is a vastly harder problem, not just because the rotation rate is smaller, but also because you do *not* have the luxury of stopping, or reversing, its rotation for purposes of comparison!

7 The Mach-Zehnder Interferometer

The Mach-Zehnder interferometer represents another topology for a two-beam interferometer; its relevance to theoretical insight, and its usefulness in optical testing, motivate its construction and use. You will see that it is another two-beam interferometer, but one in which the two beams are entirely unidirectional and non-overlapping, and capable of wide separation. The ability to pass one of its two separated beams through test media of bulky proportions helps to account for its popularity and use.

a. Setting up the simplest Mach-Zehnder interferometer

In building your first Mach-Zehnder interferometer, you will want to use a HeNe laser as source, and *two* thin-slab dielectric 50/50 beam-splitters. Why two? Because in this interferometer, beam separation and recombination are going to occur at two distinct places, in fact at diagonally opposite corners of a rectangle. Figure 7-1 below suggests the topology that you will want to try first. The beam transport from laser to input beam-splitter via two steering mirrors is standard, and the two beams emerging get sent to mirrors M_A and M_C used at 45° angle of incidence just as in the Sagnac topology. But the *fourth* corner of the interferometer is occupied by a second thin-slab beamsplitter, which further splits each of two incident beams into two, to give a total of four beams emerging from the final 'beam recombiner'. To align this interferometer is to get those output beams to overlap in pairs, both in position and direction.

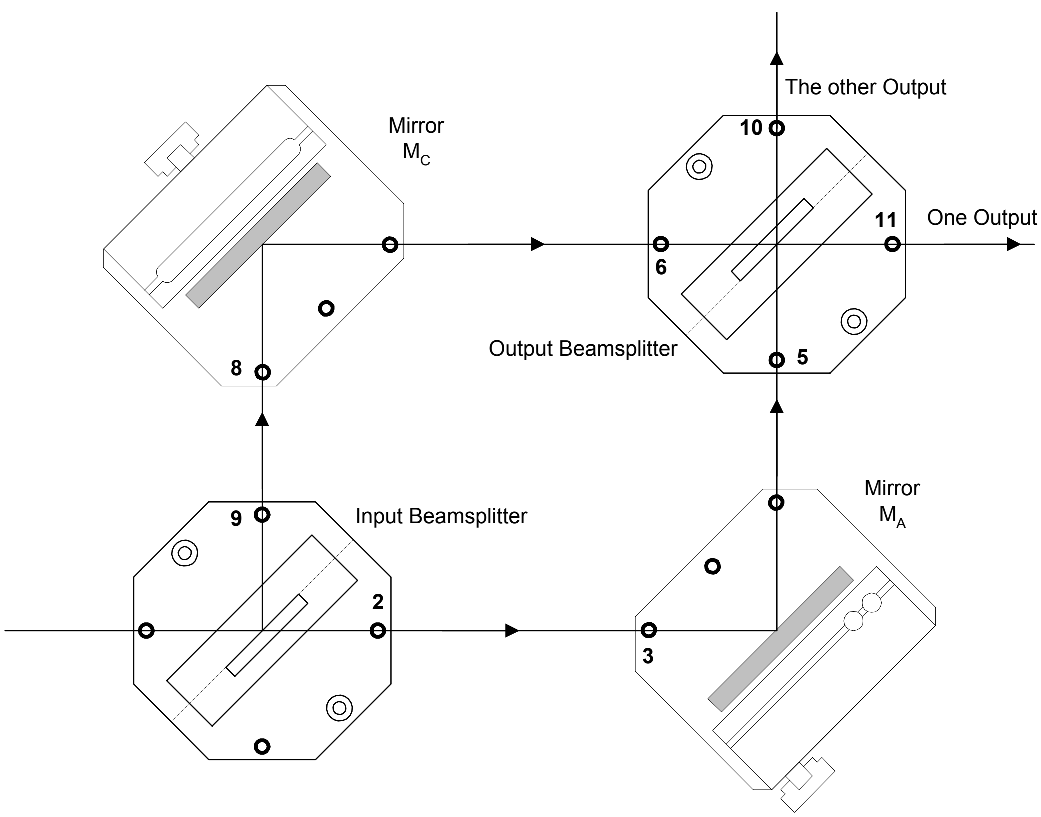

Figure 7-1: The topology of a Mach-Zehnder interferometer

Here are some hints in setting up this interferometer:

1) As in the case of previous interferometers, you want the input beamsplitter to produce a 'through' beam passing through alignment holes in beam paddles, and a 'bent' beam which is also aligned (via shimming of the beamsplitter base) to pass through beam-paddle holes. As in the case of the Sagnac interferometer, you want the bending mirrors at two corners of the rectangle to be of horizontal-hinge-line and vertical-hinge line types, to give non-redundant degrees of freedom of adjustment.

2) The fourth corner of the interferometer is occupied by another beamsplitter-holder, with the other 50/50 dielectric-film plate in place. For optimal symmetry of the interferometer, you want the working side of the glass plates to be <u>opposed</u> as shown in Figure 7-1, so that both beams pass through equivalent thicknesses of glass (this would be crucial if you were trying to build a white-light Mach-Zehnder interferometer).

3) Now consider the labeling in the figure of a set of places where you might place beam paddles for alignment. You want the 'passed' beam from the input beamsplitter to pass through paddle holes at locations #2 and #3, and you can achieve this using the adjustments of the two steering mirrors. You want the 'bent' beam from the input beamsplitter to pass through paddle holes at locations #9 and #8, and you can achieve this by 'shimming' the input beamsplitter's base in rotation and tilt. Now put the second beamsplitter base in place, but for now ignore beams emerging from it; instead, put paddles in place at locations #5 and #6 on this base, and look at the beams coming *toward* that beam-recombiner. The goal is to get the arriving beams to pass through the holes in the beam paddles, and the thumbscrew adjustments on the backs of corner mirrors M_A and M_C provide part of the solution. The other degrees of freedom you need come from 'shimming' the tilt of one corner mirror, and the rotation of the base of the other.

4) When you have the beams reaching the beam-recombiner aligned this way, you should see that the two beams reach the active surface of the beamsplitter plate in a condition of overlap-in-position. You can put beam paddles in place at locations #10 and #11 to see that this is so; you can use your hand to block the beams farther upstream as a tool for identifying which beam spots come from which path. But it is not sufficient that two beams overlap in *position* to get the sort of interference you seek; you must also have the beams share a common *direction*. By removing the paddles at locations #10 and #11, and following the beams downstream using a white paper card, you will see that (in general) they diverge in angle. The cure for this is *not* to disturb the alignment of components previously placed; rather, the cure involves the adjustment of the last optical component, the beamsplitter serving as recombiner.

5) The mechanical fixture for that beamsplitter lacks any thumbscrew adjustments at all, but you can still rotate it (by loosening its mounting screws, rotating it against the breadboard top, and retightening). The adjustment will not be so smooth and easy as using a thumbscrew, but it need not be quite perfect, as you will see. The further degree of freedom you need to get the beams to coincide in direction (best sensed by overlap on a distant viewing card) is achieved using the tilt adjustments hidden in the baseplate of the

beamsplitter mounts. See Appendix G on shimming to understand how to use this adjustment.

6) When you have achieved a much better (though perhaps still not perfect) alignment of beams coming out in a common direction, you can return viewing screens to the still-overlapping output beams just beyond locations #10 and #11. You might look closely to see any sign of fringes, perhaps closely packed together. If you see any fringes at all, you can render them more visible, lowering their spatial frequency, by going back to fine adjustments of the thumbscrews on the backs of corner mirrors M_A and M_C. This may worsen the overlap of the beam spots in position, but it's working because it's improving the overlap-in-direction of the pairs of beams.

7) Your goal is to get adequately overlapped spots that show spatial frequency of zero, at least where they are overlapping. Now the whole overlapped output spot should flicker with interferometric sensitivity; to exercise your interferometer, you could apply gentle fingertip pressure to the metal of the uprights of either beamsplitter-holder. Or, you could put the gas cell into one arm of the interferometer, or put a tiltable glass plate in either arm. [Note how much room there is to put samples in the beams -- you could imagine one beam passing through some large piece of apparatus to permit the diagnosis of its innards via interferometry.]

8) When you can control the relative phase of the beams, you might want to pay special attention to the use of the draft cover, and vibration isolation, because this interferometer, with its widely separated beams, is uniquely sensitive to the phase differences caused by air density gradients and mechanical vibrations. But now you should see each of two accessible outputs (each due to the overlap of two interfering beams), and you should notice that the two outputs are 'complementary' again, and you should understand why this has to be so.

9) If your alignment is right, you can get one of the outputs to reach down nearly to zero in intensity, and then by a mere phase change in one arm, arrange for that output to change to a maximum, with the *other* output now near zero intensity. If you could change, in a repeatable way, between these two phase conditions (and how big a phase change would you need in that one arm?), you could 'switch' the input laser beam from emerging at one place and direction, to emerging at another place and direction. There are actual fiber-optic optical beam-switches which do operate on exactly this principle; it is the *wave* nature of light that permits a mere phase shift to be translated into a directional shift.

10) Once your interferometer is aligned, you could try using a lens, or lenses, to give a slightly divergent laser beam, or a transversely expanded beam, as your input to the interferometer. Ideally you could get a 'dark beam' at one of the outputs to expand to an entire 'dark field' over a whole area, and then you'd have a tool for sensing optical phase shift over a whole field of view. In practice, this becomes very demanding on the global flatness of beamsplitters and especially the mirrors in your interferometer, and you'll find it hard to get a good uniform dark field. But if you can get close, you are in position to do some dramatic demonstrations; a favorite one is to hold the hot tip of a soldering iron vertically below one of the separated beams, and thus to 'image' (via its reduced refractive

index) the plume of hot air rising from the soldering iron. This sort of optical diagnostic is used in wind tunnels, shock tubes, and even plasma chambers as an imaging diagnostic of the sample in question.

11) One thing you can do, even without an expanded beam, is to remeasure the index of refraction of air using this interferometer. If your answer emerges with a factor-of-2 discrepancy compared to earlier work, you can rethink your theory of the interferometer, and be reminded of another difference between a Mach-Zehnder and a Michelson interferometer.

b. A polarized-light Mach-Zehnder interferometer

All of these investigations have been done using the (nominally) polarization-insensitive dielectric-film beamsplitters, and using any old input polarization from the HeNe laser. But there are investigations of particular value for provoking deep thought about light interference and photons that are best conducted by exploiting the *polarization* of light. So now you might set your HeNe laser so it's rotated 45° about its long axis in its mount, so as to deliver equal components of horizontally- and vertically-polarized light at the entrance to your interferometer. The novelty is now to use polarizing beamsplitter cubes (PBSCs) as substitutes for both beamsplitter plates, thereby devising a polarized-light Mach-Zehnder interferometer. [To make this possible, you may need to move one of the PBSCs out of its mount in the polarimetric detector unit (where it's tilted at 45°) and install it in the second of your PBSC beamsplitter uprights, mounting that in turn on a beamsplitter base. See Appendix N for the procedure -- but be especially careful, in handling a PBSC, to touch *none* of its four polished optical faces.]

12) You might start with a working Mach-Zehnder interferometer as outlined above, and first substitute a PBSC as input beamsplitter. After installing the new beamsplitter, you will need to realign (at locations #2 and #3, then at #9 and #8) as before. You might then need to re-align mirrors M_A and M_C, using beam-paddles at locations #5 and #6 as before. Finally, you can substitute a PBSC for the second, or output, beamsplitter, and again you will need to 'shim' its base, in rotation and tilt, to achieve the far-field overlap that you need.

13) But after all this work, you will get no fringes. In fact, you'll get not four, but only two, output beams, both emerging from a single face of the output PBSC, and they will not interfere. Why so? It's time to use your Polaroids to diagnose the polarization of the beams in the two separated arms of your interferometer, and then to think about what each should do upon encountering the output PBSC. [Hint: the beam that 'passes' at the first PBSC is of a polarization that will cause it to do what at the second PBSC? And the beam that is 'bent' at the first PBSC is of a polarization that will cause it to do what at the second PBSC?] Now that you've accounted for what two beams emerge, what is the polarization of these two beams? [Block one arm at a time, upstream, to try this diagnostic on the 'other' beam.] You will again have two orthogonally polarized beams, which will fail to interfere even if properly overlapped in position and direction. To get some interference fringes, put a Polaroid, oriented near 45° as required, to pass (some power from) each beam, thereby giving two beams which *can* interfere. Finally you should see some fringes, whose spatial

frequency can be adjusted, as before, using the fine-adjust tilts on the corner mirrors M_A and M_C.

You're now ready to do some very deep thinking about light and interference. You might fine-adjust the interferometer so that a few fringes are visible in your output spot, and you might confirm that the fringes disappear when you block either of the two separated beams. You might also confirm that fringes march through your output spot when you systematically vary the optical phase difference in the interferometer. Now consider a variety of choices for the orientation of the final Polaroid, which you've put into place to render the fringes visible at all. If you orient the Polaroid at 0°, then you can know through what path of the interferometer the light has come. [In fact, use beam blocking to check by which path all the relevant light is reaching your view screen.] Similarly, if you orient the Polaroid at 90°, then you can similarly show through what path of the interferometer the light has come. [Confirm this by a similar arm-blocking test.] So in either case, you have full 'which-path information' -- *but* you get no fringes. If, on the other hand, you orient the Polaroid at 45°, then you get your fringes back -- *but* you entirely lose the 'which-path information' [and can in fact confirm that energy is passing equally through both arms]. Other settings of the Polaroid will give partial which-path information, and complementary fringe visibility; in fact you could get quantitative about this, and show that there's a systematic trade-off between which-path information and fringe contrast.

Further discussion of the interpretation of all this is best left to the next section, in which the contemplation of the *photon* picture of light leads to some fascinating realizations.

8 Interlude on *Quantum Mechanics*

There are many different interference experiments, both for light and for massive particles, which use some version of a two-path geometry to produce some sort of interference effect. We bring up the relevance of quantum mechanics at this stage of your work on interferometry only because you now have an interferometer in which you have two widely separated, unidirectional, and orthogonally polarized beams in the two arms of your interferometer. You have already seen the sort of 'complementarities' that exists, at the level of classical light waves, between knowing two sorts of information:

- which-path information is knowing which, of the two arms, was the route of transmission of the energy from source to detector;
- relative-phase information is knowing the phase difference between the two arms of the interferometer, as deduced from the state of the fringe pattern.

The complications that are about to occur are *not* in the operation of the Mach-Zehnder interferometer, which blithely goes on doing what it's been doing; the complications are in *forming a mental picture of how it's doing that*. They start with the realization that your apparently continuous beam of light is in some sense a stream of photons, produced in your source, passing somehow through the interferometer, and causing individual photon events at your photodetector. With milli-watts of optical power, the photon arrival rate is so high that you cannot sense the individuality of these photon events, and in any case the photodetectors you're using don't permit detection of individual photon events. But you should know that the light intensity could be turned down far enough, and the detector could be changed to one capable of single-photon detection, and the apparatus would still show the same phenomena.

So imagine a polarized-light Mach-Zehnder interferometer, equipped with a final Polaroid between the output beamsplitter and the detector. You know that if that Polaroid's axis is oriented at 0° or 90°, then there is no interference and no two-path complications: you might have the picture that any photon detected has come from the source via an unambiguous single path, and that for this very reason there's no expectation that the photon count rate should depend on any phase difference between the two paths. The complication comes when the final Polaroid is oriented somewhere near 45°, in which case you *can* detect 'fringe signals'. To be precise, the count rate of individual photon events would display a sinusoidal dependence on the optical phase difference between the two paths. [Note this requires the detection of a whole *sample* of photons; but each such detection is of a single photon.] What's more, you can easily set the source's intensity so low that the typical time interval between photon events is much *larger* than the time a photon spends going through the interferometer. Thus your mental picture has to have photons passing through the interferometer 'one at a time'. Now the question is -- does a given photon pass through one arm, or the other?

- If it *does*, it's hard to imagine why a phase shift in the *other arm* ever makes any difference.
- If it *doesn't*, it's hard to imagine what that photon *is* doing to get through the interferometer.

There are further paradoxes to consider. Note that if you change the apparatus to place a detector 'inside the interferometer' (and there's lots of room for doing so), then every single photon detected is definitely in one arm or the other, and each photon thus detected also has a definite and unambiguous linear polarization. But since those linear polarizations for the two arms are orthogonal, it's harder still to imagine what is interfering.

There are yet more curious issues that arise when you think about photons that are parts of 'optical pulses'. You should know that it's possible to produce a burst of photons in a pulse of 1 ns, or even 1 ps, duration and so you can think of a 'bunch' of photons with longitudinal extension of just 0.3 m or 0.3 *mm* for these two cases. Now consider such pulses, but with intensity so low that there's at most one photon in a bunch. Now you can think of the three separate time intervals in which the 'bunch' has not yet entered the interferometer; when it's within the interferometer; and when it's exited the output beamsplitter. *After* that transit time of the bunch (with its one photon) through the interferometer, you could *then* finally decide on the orientation of the final Polaroid. This is an example of a 'delayed-choice' experiment, in which you seem to decide 'after the fact' whether you're going to get which-path information or phase-difference information. It takes some considerable care to form a view of what is going on that doesn't permit the highly implausible notion that you can 'change the past' by subsequent choices of what you choose to measure.

You might want to read a lovely tutorial article entitled "Quantum Erasure" by Stephen P. Walborn et al., [in *American Scientist*, vol. 91, 336-343 (2003)] about delayed-choice experiments, to see how they can be analyzed; in the process, you are certainly going to understand more of the amazing subtlety that needs to be used in discussions of the photon picture of light. You might also want to know that experiments equivalent to Mach-Zehnder interferometry can nowadays be conducted with massive particles, such as neutrons or whole atoms; see if your mental picture for an optical interferometer still works when applied to such 'honest-to goodness particles' as these.

9 Experimental Wavelength Measurement

An interferometer makes possible the direct measurement of the wavelength of light, by a technique nearly as direct as using a ruler. The measurement technique is most easily conducted in a Michelson interferometer, and it provides a direct connection between the wavelength of light and the machinist's use of ordinary length-measuring instruments. You'll see here how to measure the wavelength of HeNe and diode-laser light, using a micrometer to control and measure the translation of a mirror by small but measurable macroscopic distances.

These measurements make use of a simple property of a Michelson interferometer: a translation of one of the end mirrors (parallel to the direction of the light's travel in that arm of the interferometer) by a distance of $\lambda/2$ is just what's required to add a distance of λ to the round trip flight of light in that arm, and just enough to take you through one full fringe in the detected interference pattern. The trick is to be able to move a mirror far enough that the translation can be quantified using ordinary technology. That requires the solution of two problems:

1) translating a mirror by a millimeter is easy; translating it by 1 mm while preserving its orientation to a small fraction of a milli-radian is much harder. The TeachSpin interferometer includes the 'flexure translation stage' as an option for the baseplate of one of the end mirrors, and when suitably used, it allows a translation of an end mirror by ±1 mm from a central position, preserving over that whole distance an alignment good enough to preserve adequate fringe contrast in the detectable interference signal. You will want to consult Appendix F to learn how this translation stage works, and how to control it.

2) translating a mirror by a millimeter is easy; translating it with control and resolution on the micron (1 μm) scale is harder. The TeachSpin interferometer borrows from the world of precision machining a tool called a 'differential micrometer', which allows mechanical displacements of millimeter size to be generated, controlled, and quantified with a resolution of 0.1 μm, which is 10^{-4} mm. Appendix F also covers how this 'diff mike' works, and how to read its scale to this high a resolution.

To use these remarkable tools, set up a Michelson interferometer with nominally equal arm lengths. You might wish to have the ordinary translation stage, and its ordinary 0-1" micrometer control, in one arm of the interferometer, but you will certainly want the flexure translation stage in the other arm of the interferometer. After you line up the beams in the ordinary way, you will see two output beams overlap, but you will fail to see fringes. This is because the unrestrained flexure stage does in fact flex, and allows its moving stage, and the upright and mirror attached, to vibrate back and forth with many microns' amplitude and at a rather high frequency. The result is a fringe signal that changes too rapidly for your eye to follow it. So you'll need to add to the flexure stage a means of controlling the position of the 'stage' relative to the 'frame' of the device; this is typically achieved by using a 'pushrod' in opposition to a 'push-spring', as shown in Figure 9-1 below.

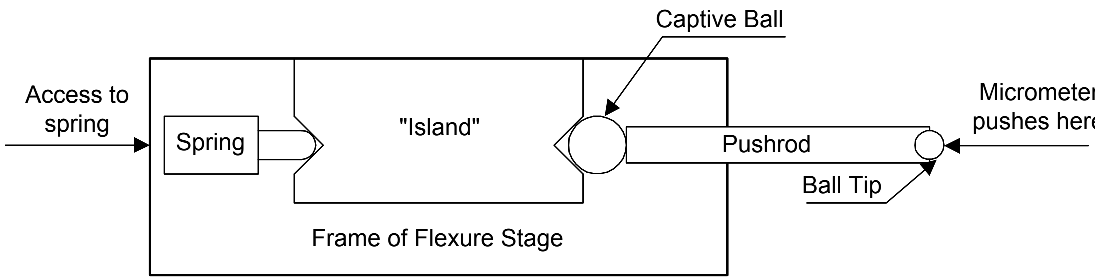

Figure 9-1: Side view of push-spring and pushrod for driving the flexure stage

The ‘spring’ is recessed inside the frame of the flexure base, and can be extracted and examined (use a 5/64" Allen wrench to engage its outer end). When you reinstall it into the flexure base, you can see its spring-loaded tip engage with the moving stage, and start to push the stage away from its central position. You should insert the spring far enough to push the stage by most of the 1 mm of motion allowed by the limit screws also visible in the flexure base.

Now the role of the pushrod is clear; bearing on the other side of the stage, it can push *back* against the spring, moving the stage back by 1 mm to its central position, and up to 1 mm beyond this, for up to 2 mm of travel. The position of the pushrod’s far end is intended to be controlled by the differential micrometer. Once a rigid pushrod is in place, under compression between a micrometer on one end, and the push back from the push-spring on the other, the free vibration of the flexure stage will come to an end. Then the fringes will stabilize, and the alignment of the interferometer can be accomplished in the usual way.

Now learn to read the setting the diff mike, which advances its non-rotating working surface by 50 μm for one full rotation of the barrel, whose smallest division on the barrel is 1.0 μm, and whose setting can be read to 0.1 μm using its Vernier scale. Finally you have a readability that approximates the tiny translations you’ll be making for single fringes.

As you learn to use the diff mike to make sub-fringe adjustments of the interferometer, you’ll notice that the mere pressure of your touch on the diff mike can distort the optical breadboard sufficiently to give fringes’ worth of deformation. The cure for this problem is to drive that micrometer in a hands-off way, using a motor-drive unit that can engage the micrometer’s barrel, and apply torque to it, without applying extraneous force to the optical breadboard. Figure 9-2 shows the motor drive in use with the diff mike; note the use of a ‘spline drive’ that allows the micrometer’s barrel to translate by 10-20 mm while the motor executes pure rotation. Note also the motor runs in one direction only, such that the micrometer’s barrel approaches closer to the motor (meanwhile letting the flexure-mounted mirror to recede from the beam-splitter). There’s a limit switch underneath the motor mount that will shut down the motor when the spline coupler comes too close.

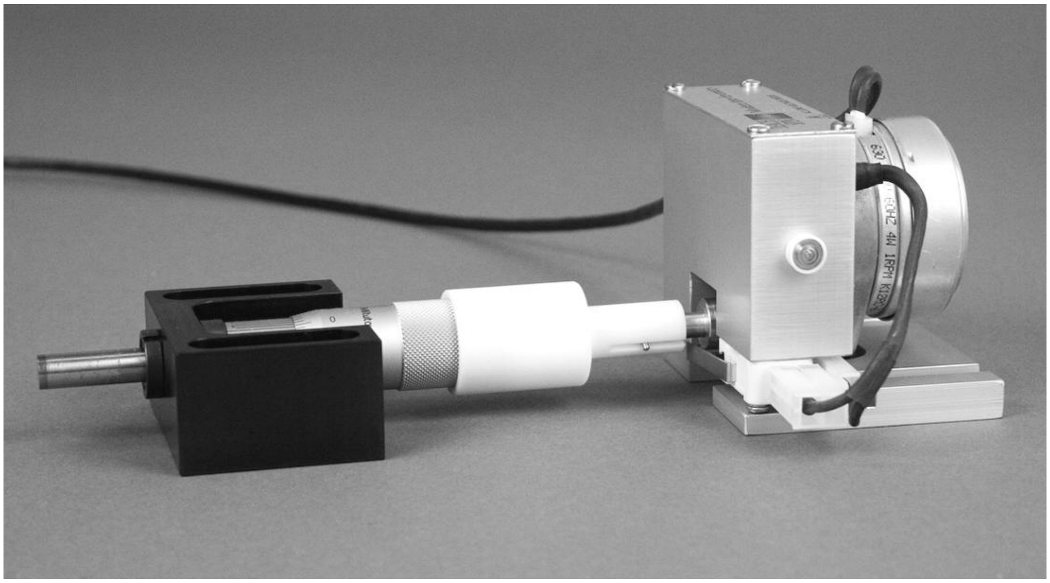

Figure 9-2: The motor drive coupled to the differential micrometer

Because the motor turns in only one direction, and a direction that 'backs out' the barrel of the micrometer, you'll need to mount it to the optical table in such a way that its spline-driving cross pin engages *shallowly* the spline cut into the white plastic adapter. To engage the pin in the slot, you'll need to hand-rotate the micrometer to match the pin's azimuth, and then position the motor until its shaft is coaxial and parallel with the shaft of the micrometer. Once you have the pin engaged, you can check that electrical drive of the motor will cause mechanical drive of the micrometer, at the motor's nominal rotation frequency of 1 revolution per minute (for 60-Hz ac drive). You'll need to watch that the drive continues to work smoothly for multiple turns, and you'll appreciate that the limit switch will eventually shut down the motor (*before* it causes the plastic sleeve to collide with the motor drive).

Once you have the motor working mechanically, you should monitor the interferometer signals that it causes. You should see fringes coming by, smoothly and at a nominally constant rate. You might first want to try counting fringes using single-channel detection, which is optically the simplest case. In this case, you might follow your detector signal through the Modern Interferometry controller box, and test the usefulness of low-pass filtering on your sinusoidal fringe signal, hoping to filter out the high-frequency noise while preserving the low-frequency sinusoid of interest. You might also check what amplitude of signal you're getting, and use an appropriate amount of hysteresis in sending your signal to a single channel of the counter. When you have found a combination that gives reliable counting, giving one electronic count for each and every optical fringe, you are ready to measure wavelengths.

The algorithm of Section 1d) will help you understand the complications of mechanical slack and backlash, and their solutions. Attempts to count about 10^2 fringes will introduce you to any complications or difficulties; naturally, attempts to count about 10^3 fringes will give you the greatest precision.

Once you have learned how to measure the wavelength of light in concept, you might want to use the quadrature-interferometry scheme to do actual measurements. This technique has several advantages. First of all, you get four rather than one count per fringe, so your measurement resolution goes up four-fold. Second, you can be sure that transients in the rate of motor driving the diff mike don't matter, since your system now counts reversibly. Thirdly, your system can now count fringes faster; you might even want to revert to *hand*-driving the diff mike (or even the 0-1" micrometer on the other translation stage) through a desired range.

After you've established your level of repeatability, and your best-measured value for the HeNe laser's wavelength, you might reflect on the potential for systematic errors in an interferometer like this. You might think about 'cosine errors' and 'Abbe errors', and reflect that the actual motion of the face of your micrometer's shaft, and the face of your flexure's mirror, could differ for several subtle reasons.

When you're ready, you can substitute the diode laser for your HeNe laser, and get a wavelength to measure for which there is no 'book value' to look up. Your red diode laser belongs to the nominally 650-nm class of diode-laser technology, but it might actually operate anywhere between 645 and 660 nm. The output might fail to be of a single frequency, so that you might suffer from reduced coherence length in your interferometric measurement. Furthermore, the wavelength of the laser is somewhat temperature-dependent; you might even try to verify this, by using the thermoelectric modules attached to its mount to vary its temperature deliberately. Finally, you might look up 'wavemeters', to learn how they measure the wavelength of an incident laser beam.

After all the details of wavelength-measurement exercises for HeNe and diode lasers, you might reflect a bit on the kind of capability you have demonstrated. What you have accomplished is to establish a *ratio* between the wavelength of light and the calibration marks on a length-measuring instrument like a micrometer. From a fundamental point of view, the *significance* of this ratio changed when the SI definition of the meter was revised in 1960 and again in 1983. Previous to 1960, when there really was a 'meter bar' preserved in Paris, a micrometer was calibrated by reference to the meter bar, and the wavelength of light really was an unknown which could be measured via an experiment of the sort you've performed. But subsequent to 1983, the meter's definition has changed, and there no longer *is* a 'standard meter bar' of any authority. Instead, an SI meter is *defined* to be the distance light travels, in vacuum, in a time of (1/299,792,458) of a second. With the speed of light thereby defined, the wavelength of light of a given frequency can no longer be measured; instead, it will be fixed by $\lambda = c/f$, where f is the frequency of the light. For a laser like a HeNe laser tied to an atomic transition, that frequency is fixed and can be measured; your HeNe has a frequency is near 473,613 GHz, and the (vacuum) wavelength is therefore 632.990 nm.

So what *are* you accomplishing, post-1983, when you go through the series of tabletop manipulations that used to be called 'measuring the wavelength'? You're still measuring the ratio between the wavelength of light and the calibration marks on a length-measuring instrument; but now the unknown has changed. Formerly the meter was defined, so via that ratio you were measuring the wavelength; now, the wavelength is no longer in question, *so you are 'realizing the meter' on your tabletop*, and establishing the calibration of your micrometer in relation to the meter thereby realized. In fact *all* standards of lengths have to be established, under the current definition of the meter, by some process analogous to this; and in fact all micrometers have their calibration ultimately traced back to some process akin to this.

10 Experimental Thickness Measurement -- An optical 'keepsake'

You've seen (in Section 9) the usefulness of interferometry in measuring the wavelength of light, but what about the more concrete measurement of the length of a material object? Here's an exercise that uses the bi-directional counting scheme of quadrature interferometry (section 3) and applies it to the measurement of the thickness of actual experimental samples.

The motivation for this experiment comes from length metrology. If you've used (or heard of) a 'feeler gauge', you can imagine a metal sample with two plane-parallel and smoothly polished surfaces, and you can imagine that a manufacturer can claim (for example) that the gauge's thickness is 0.0250" = 0.6350 mm. Beyond the thicknesses given by feeler gauges are the lengths of 'gauge blocks', again made with highly polished parallel working surfaces, and again manufactured to serve as standards of length.

What this experiment does is to show you how a feeler gauge or a gauge block gets its *authority*; that is, by what process the thickness or length stamped on it derives its connection to the SI definition of the meter. In fact, in this experiment, you'll measure the thickness of a feeler-gauge sample directly by interferometry.

In this illustration of the method, you'll be limited to measuring thicknesses in the 0-2 mm range, only because the flexure stage is limited to such a range. You'll be able to measure any thin sample that you can introduce between the differential micrometer's polished face and the invar pushrod's 3/16" ball-bearing tip. And you'll be able to measure any thickness in this 0-2000 μm range to a least-count resolution of about $\lambda/8 = 0.079$ μm = 79 nm.

Figure 10-1 shows the basis of the method.

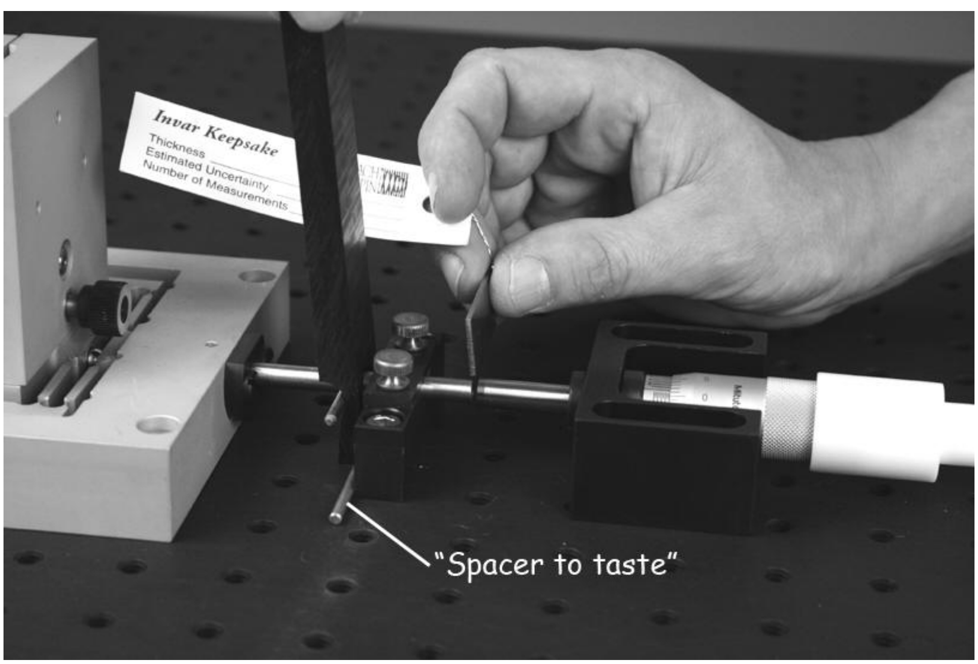

Figure 10-1: The hand lever in place, ready to open a gap between the pushrod's ball-tip and the micrometer's working face

With the diff mike's front face, via the pushrod, determining the position of one end mirror of a quadrature Michelson interferometer, you need a method for controllably opening a small *gap* between the mike's front face and pushrod's ball tip. Into this gap, a sample can be introduced. You need to think of a sequence, in which

- you start with that gap at zero, the ball tip contacting the mike face;
- then you more than adequately open a gap, from zero to (say) 1.5 mm;
- next you slide a sample (say, 1.2 mm thick) into the gap;
- finally, you allow the gap to close, with the sample now pinched between the ball tip and the mike's face.

Can you see why bi-directional counting is essential to this method of measurement? Can you see that the micrometer is *not* being used to measure anything?

In order that the counting method should work, the whole process of opening and closing the gap needs to be accomplished without loss of any counts. This requires smooth enough motions at controlled speeds. Surprisingly, with a little practice this can be accomplished *by hand*, using the black plastic 'lever arm' shown in Figure 10-1. If you have a quadrature interferometer set up and working, you can drop this lever into the space between the pushrod's brass crossbar and the V-block, and use the lever (by pushing its top to the left) to move the pushrod to the left, and thereby open up the desired gap at its right end.

Before you think about actually measuring a sample's thickness, you should try your hand at measuring a 'thickness of zero' by this method. That is to say, you should start with

- a gap of zero, pushrod's tip contacting mike's face, and
- a count of zero, controller's counter reset to 00000

and then open up a gap of a millimeter or so. (You'll be using the lever's mechanical advantage to push against the flexure stage's return spring, and you'll feel the force required.) Now the question is: when you relax your hand and allow the gap to close, will the count return (from a temporary and wildly fluctuating reading of many *thousands*) to a reading of zero?

You can appreciate that the count can return to zero *only* if the pushrod's ball tip makes a repeatable contact with the very same spot on the diff mike's polished face. That's why there's a sleeve in the flexure, and a V-block on the breadboard, to ensure that the pushrod moves only along a straight line. You will perhaps intuit, or learn, that you don't want the pushrod to rotate during this one-dimensional translation; the pushrod's crossbar is intended to interact with the lever arm in such a way as to keep that brass crossbar level. Finally, you might see how close you are to the state-of-the-art in repeatable metal-to-metal contact by estimating a few relevant distances, and expressing them all in the same set of units:

- what is the motion of the pushrod that gives you one single count on your counting system? (Remember you get four counts per full fringe in the quadrature mode.)
- what is the surface roughness of polished steel surfaces, like your ball-bearing tip or the micrometer's front face?
- what is the 'sphericity' of a ball bearing, and how is it relevant here?
- how much does the pushrod's length compress under the compressive load that it undergoes?

- how much does a ball bearing flatten when it's pressed under load against a plane surface?

After you've learned the technique, and established the reliability, and level of repeatability, for 'measuring a thickness of zero', you can try measuring some actual samples. If you practice using a glass microscope slide, you can imagine making a whole series of readings, alternating between 'sample out' and 'sample in' conditions, while *not* resetting the counter during the series. You might even come to find that microscope slides do not *have* a well-defined thickness, at the level to which you can measure.

So you might try some samples that *do* have accurately plane-parallel faces. You could try stainless-steel foil; you could try actual steel feeler gauges; or you could try one of TeachSpin's invar 'keepsake' samples, whose surfaces have been lapped to be mirror-smooth and accurately plane-parallel (and can you detect *departures* from perfect parallelism?). With care, you should be able to assign a precise, and accurate, thickness to your keepsake sample, and (hardest of all) even assign to your value a defensible uncertainty of measurement.

11 Measuring Indices of Refraction

The passage of 'fringes' in a interferometer depends on changes in the phase shift in one beam of the interferometer relative to the other. One full cycle of some detectable quantity goes by when a 2π-radian phase shift occurs due to some change of an independent variable. In this section you'll see that the measurement of phase shift via a fringe count can be the basis of a method for measuring an important property of materials, the index of refraction. This section is written assuming that it's a Michelson interferometer that's being used to make the measurement; the same techniques can be used in Sagnac or Mach-Zehnder interferometers, provided provision is made in the theoretical treatment for the fact that these latter devices involve a one-way (rather than two-way) passage of the light through a sample.

a. The index of refraction of gases

TeachSpin provides a 'gas cell', an optical sample chamber with two transparent end windows, which can be placed into the beam in one arm of an interferometer. The idea is that as the pressure of gas in the cell varies from near-vacuum to near-atmospheric, the density and the index of refraction of the gas will vary, and that in turn will change the phase shift in that arm, and cause a succession of fringes to appear at the interferometer's output.

It's worth recalling that the index of refraction of a material, n, is defined by $n = c/v_{ph}$, where the phase velocity of the wave is in turn given by $v_{ph} = \omega/k$. The angular frequency ω is connected to more familiar quantities via

$$\frac{\omega}{2\pi} = f = \frac{c}{\lambda_{vac}} \quad ,$$

so these combine to give the wave number k as

$$k = \frac{2\pi}{\lambda_{vac}} n \quad .$$

The relevance of the quantity k is that a (complex) wave propagates in space according to

$$\exp i(k\,x - \omega\,t) = \exp\left(i \frac{2\pi}{\lambda_{vac}} n\,x\right)\; e^{-i\,\omega\,t} \quad .$$

Hence a sample chamber of fixed (internal) length L, under a change of index of refraction from $n = 1$ (in the vacuum condition) to a given n-value, causes a phase shift for one-way travel of light given by

$$\Delta\phi = \frac{2\pi}{\lambda_{vac}} \Delta n\, L = \frac{2\pi}{\lambda_{vac}} (n - 1)\, L \quad .$$

In a Michelson interferometer with two-way passage of the light, the phase shift is double this, and the fringe count M between 'vacuum' and 'full' conditions is given by

$$M = \frac{\Delta\phi}{2\pi} = \frac{1}{2\pi}\, 2\, \frac{2\pi}{\lambda_{vac}} (n - 1)\, L = \frac{n - 1}{\lambda_{vac}} (2\, L) \quad .$$

The TeachSpin gas cell has internal length L = (100.0 ± 0.1) mm, and when used with red HeNe laser radiation of vacuum wavelength λ_{vac} = 632.990 nm, it will give about M = 85 fringes even for the small change in index of refraction of air from 1.0000 to about 1.0003 which occurs between vacuum and 1-atmosphere pressure. So if you've been used to assigning "1" to the index of refraction of air, now you'll be able to quantify departures from n=1 with a precision of a part per million!

There's some good physics in the index of refraction of a gas, in its dependence on the pressure and temperature of the gas, and the polarizability of the molecules composing it. The behavior of gases and even liquids is very well described by the Lorenz-Lorentz law, and this in turn allows the measurement of the index of refraction of a gas at a given temperature and pressure to be corrected back to the values expected for STP [standard temperature and pressure] (which in the optical community is chosen to be 15 °C and 760 Torr).

Depending on what gases you have available for study, you might try measuring the index of refraction of air, nitrogen, oxygen, helium, argon, or hydrogen. You might also understand how the electric polarization of a gas *mixture* is related to the polarizability and the number density of the molecules composing it, and learn how a measurement of the index of refraction of a mixed sample can be a simple surrogate for measuring the mixing ratio of its constituents. If you look up quantum-mechanical calculations for the polarizability of the helium atom computed from first principles, you'll be able to compare measured and predicted values for its refractive index.

The TeachSpin apparatus makes the pressure dependence of index of refraction easy to study, since it provides a piezo-resistive differential-pressure transducer that translates gas-pressure difference between its two ports into a real-time analog voltage. It can be used to measure cell pressure relative to ambient atmosphere, or (given a suitable forepump) relative to vacuum. The nominal calibration constant of the transducer is

$$\frac{\delta V_{out}}{\delta(\Delta p)} = 10.\ \frac{\text{Volts}}{\text{atmosphere}}\ ,$$

but the actual value can be checked against some local pressure standard such as a U-tube mercury manometer or a Bourdon gauge. Even the local weather report, or airport's value of present atmospheric pressure can be used, provided you recall that these are *not* actual measured surface values of air pressure, but rather ones extrapolated down to sea level.

Your TeachSpin equipment includes a gas manifold, described in Appendix S, with valves that allow the controlled addition, or removal, of gas from your sample cell, and also the establishment of vacuum conditions on one side of your differential-pressure sensor if you wish. There's also a port to permit the controlled venting of the line to your forepump to ambient pressure. You'll want to use the flexible plastic hoses provided to permit the manifold to be mounted *off* of your optical breadboard, so that mechanical manipulations of the valves don't cause irrelevant deformations of your interferometer.

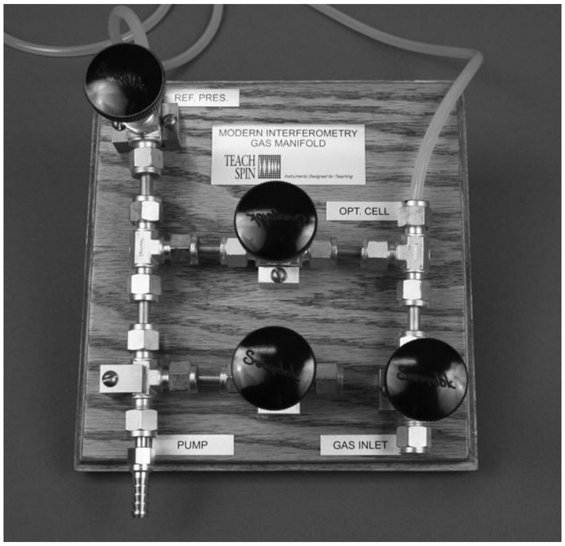

Figure 11-1: The Gas manifold

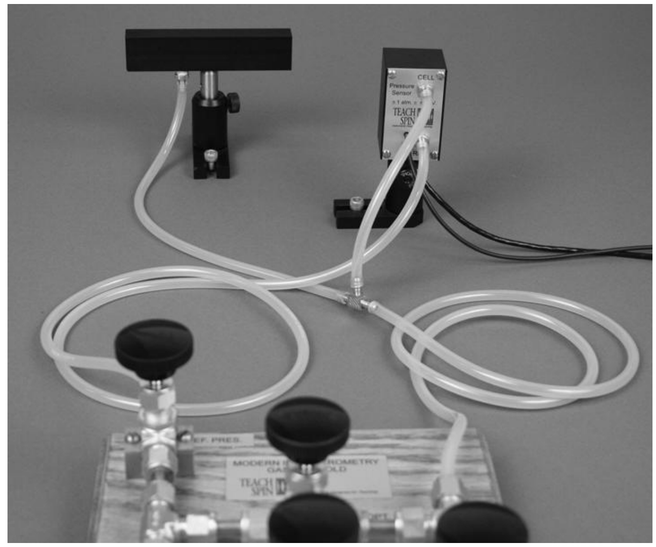

Figure 11-2: Cell and transducer interconnected

Finally it's worth remembering that nominal atmospheric pressure is 101,325 Pascals ≡ 760 Torr, and that a forepump of modest performance can be expected to yield, for present purposes, a very good approximation to a perfect vacuum. The limiting pressure of a forepump might be 10 to 50 mTorr, which is less than 10^{-4} of an atmosphere, and in any case these low pressures can be measured absolutely, if not very accurately, with a thermocouple gauge.

b. The index of refraction of slabs

Your measurement method for index of refraction of gases depends on the ability to count fringes while changing the density of a sample continuously over a range, and this method of course fails to work for solid samples. Interferometry presents another method for measuring the index of refraction of solids, provided only that they are available in the form of a parallel-sided transparent slab, which can be introduced into one arm of an interferometer. Again, you can't count the fringes that go by while you're sliding the sample sideways into the beam, so you need a cleverer way to make the sample's effect change continuously. The technique that's often used is to start with a sample already in the beam, but with its faces perpendicular to the light beam, and then to count fringes as the sample is rotated by a known amount away from perpendicularity.

The effect of a plane-parallel slab on the phase of a beam of light is complicated by the fact that there are two geometric effects going on simultaneously; one is a phase change due to index of refraction, and the other is a geometric effect due to refraction. The calculation of the phase shift expected for a slab of thickness T and index n, when tilted by angle θ relative to the face-on condition, is best computed by comparison to a mythical ray of light that misses the sample altogether. The geometry is shown in the figure below:

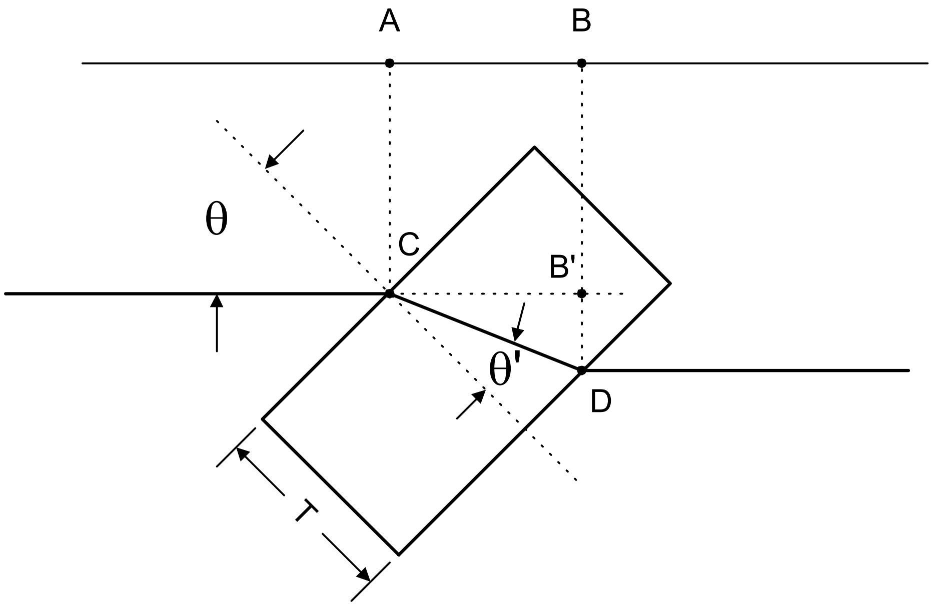

Figure 11-3: The geometry of refraction through a plane-parallel slab of thickness T, tilted by angle θ away from face-on to a beam of light

Recall that the wave number k gives phase accumulation per unit distance, so in physical distance L, the phase accumulation is given by

$$\Delta\phi = k\,L = \frac{2\pi\ (\mathrm{rad})}{\lambda_{vac}}\, n\, L \ ,$$

and that net phase accumulated is the product of this rate times the distance traversed. Hence in the diagram, the difference in phase accumulation can be written as

$$\delta(\Delta\phi) = \frac{2\pi\ (\mathrm{rad})}{\lambda_{vac}} (n \cdot \overline{CD} - 1 \cdot \overline{AB}) \ .$$

Geometry and Snel's Law give the distances involved as

$$\overline{CD} = \frac{T}{\cos\ \theta'} \ \text{ and } \overline{AB} = \overline{CB'} = \overline{CD} \cos(\theta - \theta'), \text{ where } \sin\theta = n \sin\theta' \ ,$$

so finally we can write, for the phase shift due to a slab at angle θ, relative to the same slab face-on, the expression

$$\phi(\theta) = \frac{2\pi}{\lambda_{vac}}\, T \left\{\frac{n - \cos(\theta\text{-}\theta')}{\cos\ \theta'} - \frac{n\text{-}1}{1}\right\} \ .$$

This is a messy function of the index of refraction n, the more so when the internal angle θ' is written in terms of the measured external angle θ. Using Snel's relation 1.sin θ = n.sin θ' for small values of the angles, and expanding in powers of the angle θ, one can derive the approximate expression

$$\phi(\theta) = \frac{2\pi}{\lambda_{vac}}\, T \left\{\frac{n\text{-}1}{2n}\, \theta^2 + O(\theta^4)\right\} \ ,$$

which is enough to show that the expected fringe-count signal will be approximately *quadratic* in angle θ. For a two-pass encounter with a typical glass slab having n = 1.5 and T = 1 mm, we find that the fringe count M is given by

$$M = 2\frac{\Delta\phi}{2\pi} \approx 2\frac{T}{\lambda_{vac}}\frac{n-1}{2n}\theta^2 \ ,$$

and reaches over 5 full fringes already at θ = 0.1 rad = 5.73°.

The complicated phase-shift function above thereby predicts a fringe count M as a function of θ, with a functional shape depending on the index of refraction n of the slab. The value of index n can be extracted by comparing the $M(\theta)$ data with functional shapes computed for a set of candidate n-values. Alternatively, the complicated functional dependence of fringe-count on index n can actually be inverted to give

$$n = \frac{\alpha^2 + 2(1 - \cos\theta)(1 - \alpha)}{2(1 - \cos\theta - \alpha)} \ ,$$

where $\alpha = M(\theta)\lambda_{vac}/2T$ is directly related to the observed fringe count.

Your TeachSpin equipment includes some slabs of glass and quartz that are mounted on posts in such a way that they can be introduced into a light beam and then rotated about a vertical axis by a controllable amount. There is a holder for a single slab of glass (nominal thickness T = 1.0 mm), another holding a pair of such glass slabs, and a third holding a single slab of quartz (nominal thickness T = 0.50 mm). You also have a compensator plate of 1" diameter and 2.0 mm thickness. Any of these devices can be mounted into the miniature rotation stage, which allows them to be rotated freely, or dialed through a modest range of angle under carefully controlled angular translations. So you can measure the refractive index of any of them, and check for its dependence on polarization. [Appendix P describes the structure and use of this rotation stage.]

Hence the ability to count fringes as a function of angular displacement can be used to deduce index of refraction; alternatively, these slabs and their rotator can be used to make small and very accurately controllable phase shifts in an interferometer.

12 Detecting thermal expansion

You have come to appreciate that you have, in the Michelson interferometer, a way to detect very small changes in the position of an optical element.

If you've done experiments in the direct measurement of laser wavelengths in section 9, then you have reliably created (and detected) the translation of a Michelson interferometer's end mirror by displacements on the scale of *microns*. This is a sensitivity which makes lots of other physical effects detectable, and here you'll apply it to detecting the changes in length of a sample that occur with temperature changes.

The only thing you need that's new is a sample made of the material of interest, equipped with a heater (so you can change its temperature) and a thermometer (so you can measure its temperature). Your TeachSpin kit comes with one sample (of copper) all prepared, and with some other materials ready for you to make into useful samples. The materials provided are listed here, and may be identified by their differing nominal lengths:

material	length	composition
Aluminum	60 mm	6061-T6511 alloy
Copper	65 mm	110-alloy, >99.9% high-conductivity copper
Steel	70 mm	(unhardened) steel drill rod, ASTM A681
Alumina	75 mm	sintered Al_2O_3, >94% of crystal density
Invar	80 mm	36-alloy; 36% Ni, ≈63% Fe

You can use the constantan heater wire provided, and the type-K (chromel-alumel) thermocouples included, to make up a working sample, by following the instructions that are found in Appendix Q. The HH11A readout unit provides a direct reading of temperature of your sample. You might want to look up 'constantan' to understand at least two motivations for its use as heater wire; you might want to perform a one- or two-point calibration check on your temperature-measurement system.

The thermal-expansion samples have flat planar ends. Here's a way to mount them in such that their length directly controls the position of your interferometer's end mirror:

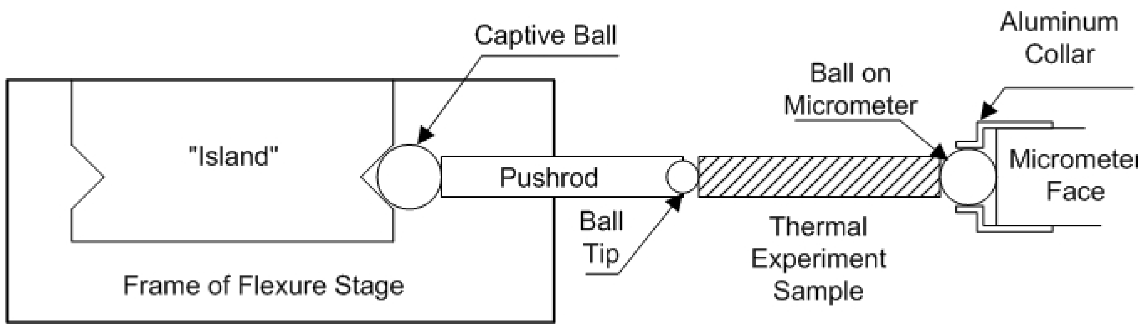

Figure 12-1: The mechanical train for a thermal-expansion sample to control the stage's position

In this diagram, note that the invar pushrod has its exposed end terminated by a 3/16"-diameter stainless-steel ball, and note also the use of a 1/4"-diameter stainless-steel ball held by an aluminum sleeve onto the face of your diff mike. These balls define unique points of contact with the flat ends of your thermal-expansion samples; the motivation for stainless steel is to minimize thermal coupling of the sample with other objects in the environment.

Once you've confirmed, by tiny adjustment of the micrometer, that motions of the sample are showing up in your interferometric signal, you are ready to park the micrometer setting, and move the end mirror via the thermal expansion of the sample. You might be using the regular, or the quadrature, Michelson interferometer, and your goal is to measure the 'coefficient of thermal expansion', given by

$$\alpha = \frac{1}{L}\frac{\Delta L}{\Delta T} = \frac{\Delta L/L}{\Delta T} \quad ,$$

with units typically written as ppm/K, or part-per-million per Kelvin. You will find that with sample lengths of order 0.1 m, and temperature changes of order 10 K, you get *lots* of fringes' worth of signal. You can now appreciate that temperature expansion of materials is a potentially serious cause of instability in interferometry in general, and you can now see in retrospect why temperature stability of the optical breadboard and the various mounts of your interferometer is important. You will also appreciate the revolution in length metrology that resulted when nickel-iron alloys of anomalously low coefficient of expansion were discovered, and why they are still used in various applications of optical metrology.

13 Detecting magnetostriction

It was James Prescott Joule, of the unit-of-energy fame, who discovered way back in 1842 that ferromagnetic materials actually change their dimensions when they undergo the process of magnetization. This process is called ‘magnetostriction’, meaning ‘shrinkage due to magnetism’. The dimensional changes are not large, so interferometry is one convenient technique for detecting them, but they reveal enough about the process of magnetization to be well worth studying. The existence of rather small dimensional changes in polycrystalline ordinary metals like iron, nickel, and cobalt has also prompted extensive research into magnetostriction in single-crystal samples, and in special alloys, and has led to the discovery of ‘giant magnetostriction’ in alloys such as Terfenol-D.

Your TeachSpin kit includes some samples and the apparatus necessary for the quantitative detection of magnetostriction. Here you’ll really want to use the quadrature-Michelson geometry, using the change in length of a long thin sample to move one end mirror in the interferometer. This time, you’ll be trying to keep the temperature of the sample constant, but to change the magnetic field in which (part of) the sample is immersed.

Your TeachSpin kit includes three rods for studying magnetostriction; each is of nominal 1/4"-diameter, and their nominal lengths and compositions are listed below.

material	length	composition
Nickel	295 mm	Alloy-200 commercially pure (>99%) nickel
Steel	300 mm	(unhardened) steel drill rod, ASTM A681
Copper	305 mm	110-alloy >99.9% high-conductivity copper

These samples, or any other of similar shape and size, can be fit into the mechanical train illustrated below. With the micrometer set to a convenient fixed location, the sample rod will be controlling the position of the end mirror attached to the flexure translation stage. You’ll note that the sample rods have flat planar ends, and that they make contact with spherical balls at both their ends to provide well-defined mechanical contacts. There’s a 9/32"-diameter ball held captive inside your flexure mount, and you can use a 1/4"-diameter ball held to the face of the diff mike by an aluminum sleeve.

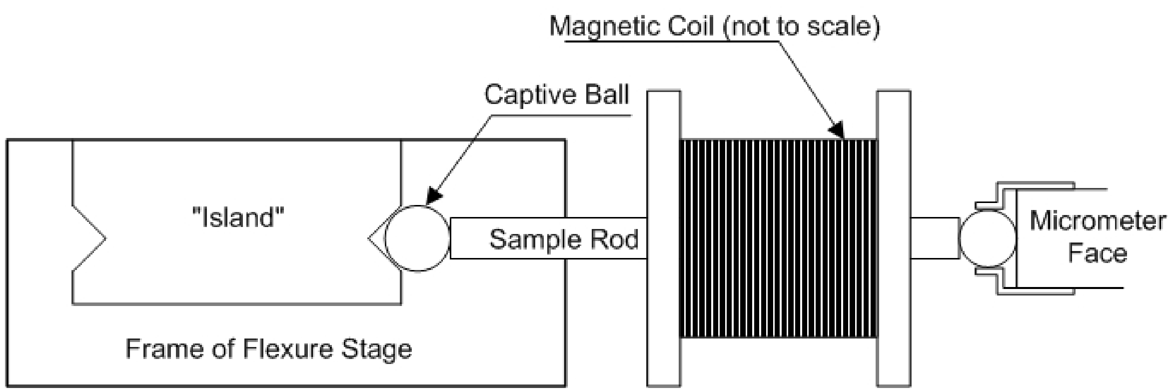

Figure 13-1: The mechanical train for a magnetostriction sample to control the stage’s position

With the mechanical train holding the sample rod on a centerline 0.70" above the tabletop, there is now room to surround a 0.161-m long portion of the sample with a solenoid, and thereby to subject it to a magnetic field. [There's also room inside the solenoid to put in place some melamine plastic foam pieces around the sample rod.] The current in the solenoid is under your direct control, and that current in turn directly controls the 'magnetic intensity' field H, according to the result

$$H(t) = n\, i(t).$$

Here i(t) is the current you control, and n is the turn-density of the solenoid, the number of turns per unit length. The multilayer solenoid has a total of 2400 ± 60 turns of #22 AWG copper (wound in 10 layers) between its end plates, and tolerates continuous currents of up to 3 A, permitting you to reach fields up to $H = 43.8 \times 10^3$ A/m.

Why use this oddball field H? Recall that in the presence of magnetizable material, the more familiar field B is given by

$$B = \mu_0(H + M)\ ,$$

where $\mu_0 \equiv 4\pi \times 10^{-7}$ T.m/A, and where M is the magnetization (the volume density of magnetic dipole moment) of the sample. In ferromagnetic materials, M is a complicated function of H, *and* of the material's past history, so that B is *not* directly under your control. It's also the case that M can be much larger than H; your solenoid's field H in some sense 'leverages' the magnetic moments already present, to yield a field inside the sample much larger than just $B = \mu_0 H$.

So with H as independent variable, your goal is to study ΔL, the changes in length of the sample. You might want to start with a copper sample, to serve as the 'control group' in your experiment, and to learn what ΔL effects can occur that have nothing to do with magnetism. Then you might want to switch to the tool-steel sample, to see if you can detect effects that definitely are related to magnetism. Recall that your quadrature Michelson interferometer can be used with the XY-'scope display, in which a quarter-turn of rotation on your display locus corresponds to a round-trip change of a quarter-wavelength, and hence a one-way mirror motion, controlled by your magnetostrictive sample, of just λ/8. It's conventional to quantify magnetostriction in microstrain units, where the quotient $(\Delta L)/L_{sample}$ is officially the 'strain' of the sample, and where $(\Delta L)/L_{sample} = 1 \times 10^{-6}$ defines one microstrain. You might also think about what is the effective length L of your sample. Finally, you might use the quadrature technique to be sure you can tell when your sample is growing in length, and when it's shrinking; you can always use a very gentle fingertip pressure, in a definite direction, on the mirror mount, to see the signal that arises from a translation of definitely known *sign*.

The full glories of magnetostriction are best observed using the polycrystalline nickel sample. Here you should find signals large enough to motivate a thorough quantitative study, but at first you will seem to have a hard time getting reproducible data. The reason is that you are discovering the *history*-dependent length of your sample; the interferometer's readout of the present length will depend not only on the present setting of your solenoid current, but also on its past history. In particular, you will get a different result at (say) i = 0.5 A if you're approaching that current from above, than if you're approaching it from below. So the most reproducible, and understandable, data will be taken with the sample starting in (say) a fully magnetized condition, and then taking data at a variety of points during the monotonic decrease in current, and another set of data at a variety of points during the monotonic increase in current, back to a starting point. The most symmetrical sets of data will emerge if your scans over current can cover a range symmetrically disposed around zero current, so your independent variable H covers the full range from $-H_{max}$ to $+ H_{max}$. Now you can see a full magnetic cycle, and perhaps intuit what's going on at those remarkable places where a monotonic increase in magnetic field H nevertheless causes a sudden *reversal* in the length change ΔL. You can also gain a fresh appreciation for the usefulness of quadrature interferometry, since such reversals could be so very easily missed in ordinary Michelson interferometry.

In a more complete study of magnetostriction, you would want to explore a range of field strengths up to the point of magnetic saturation, or a range of temperatures up to the Curie point; or, you might include the simultaneous measurement of the magnetostriction and of the magnetization M of the sample. To see the state of the art, you might want to look into 'smart materials', and find some of the uses already being made of magnetostrictive alloys as mechanical actuators. You might also compare the size of the length changes you have achieved with those attainable in the record-setting alloys that have been discovered since Joule's time.

14 Applications of piezoelectricity -- In pursuit of the nanometer

Piezoelectricity is the production of electric charge separation due to pressure exerted on a material, and was discovered by Pierre and Jacques Curie around 1880. It is a remarkable fact that the symmetrical compression of a non-conducting slab of material can result in the appearance of positive charge on one face, and negative charge on the other, because a symmetry argument would suggest that there is 'no sufficient reason' distinguishing one face from the other. In fact, piezoelectricity can only occur in crystalline materials whose unit cells *lack* the property of inversion symmetry, and therefore its macroscopic manifestation is a window into the microscopic ordering of atoms in a crystal.

But the use of piezoelectricity in optics is more pedestrian, and also reciprocal; typically, one applies a potential difference to a sample of piezoelectric material, and exploits the mechanical deformation that occurs in response ('converse piezoelectricity'). Often, the deformation is used as a mechanical translator, a direct conversion of electrical signal into mechanical motion. For typical bulk materials, the size of the deformation is really tiny, with only micrometers of deformation even for kilovolts of potential difference. But deformations of this size are useful in optics, and larger deformations can be achieved using artfully stacked slabs of material with layered electrodes in place to multiply the effect of a given potential difference.

Your TeachSpin suite of experiments includes one commercial realization of a 'piezoelectric translator' that you can investigate in detail. It's a rather small stack of crystals and electrodes, with the attractive property that its length (of under a cm) can be made to change (over a range of order 3 μm) in response to a potential difference in the range 0 - 100 Volts. The commercial device is packaged with a metal fixture that allows you to put the 'piezo stack' into the mechanical train that fixes the position of the flexure stage of the TeachSpin flexure base for end mirrors. By this means, you can translate the end-mirror over this range, large enough to give dramatic effects in the interferometric context.

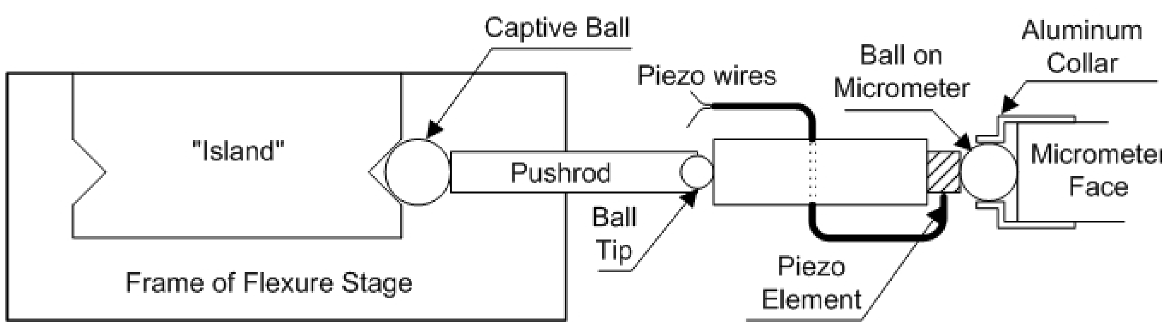

Figure 14-1: The mechanical train for a piezoelectric actuator to control the stage's position

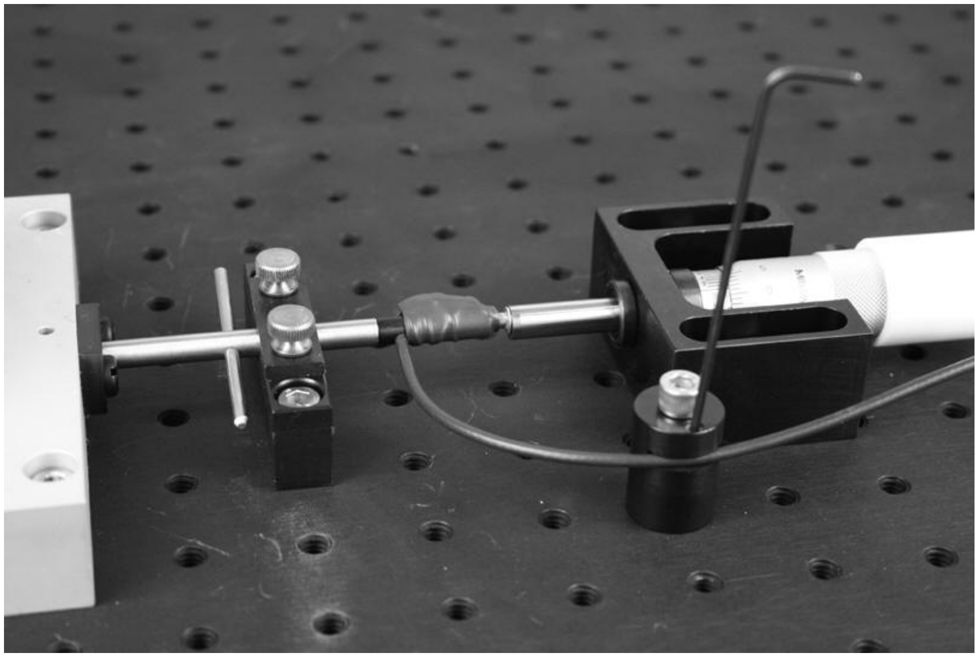

Figure 14-2: The piezo stack and the clamp for its connecting wire

Two kinds of safety notes: any piezo translator will work reciprocally as a piezo generator, which means that quite modest compressive pressures applied to the piezo stack can result in the appearance of kilovolt potential differences on the open leads connected to its electrodes. [Think of the operation of kitchen ‘sparkers’ for ignition of flames.] There’s no steady current thereby generated, but the discharge of the pressure-induced charge separation through your fingers would be a nasty surprise. Second safety note: the asymmetric (‘poled’) crystals in your piezo stack also have an asymmetric range of safe operation for potential difference. Relative to the potential of the outer conductor, you are to keep the potential of the inner conductor confined to the *un*symmetric range -10 V to +110 V.

It would be handy to have a dc power supply, adjustable through the range 0 - 100 Volts, to actuate the piezo stack, but you can easily detect effects even in the smaller range 0 - 15 Volts. The power supply needs to supply only trifling currents, and needs to supply them only during changes in voltage. A piezoelectric stack behaves electrically as a capacitor, with charge on either plate given in terms of the potential difference ΔV by

$$q = \pm C\,\Delta V \quad ;$$

hence the current flowing in the wires to the device is given by

$$i = \frac{dq}{dt} = \pm C\,\frac{d(\Delta V)}{dt} \quad .$$

So a low-current supply will only limit the rate at which you can change ΔV; mere μA supply capability will suffice if you’re not in a hurry.

Your goal is to make an interferometric measurement of the extension of the piezo stack as a function of ΔV. You should be able to test whether a positive potential applied to the central conductor of the coaxial cable causes the device to expand, or to contract, via comparison with the known direction of motion caused by your micrometer. You should be able to find the maximal extension you can achieve at the maximum ΔV you choose to apply. But you can get more quantitative than that, by noting the detailed dependence of length changes ΔL on your independent variable, ΔV. If you start at $\Delta V = 0$ and increase the potential difference monotonically to a maximum, taking data points on the way up, you will get one plot of the function $\Delta L(\Delta V)$; but you should also take another set of data on the way *down* from your maximum potential difference, and get another plot of $\Delta L(\Delta V)$. These plots may fail to coincide; you have discovered *hysteresis*, in the sense that the sample's extension ΔL depends not only on the present value of ΔV but also on its past *history*. The technological use of piezo translators may need to take this effect into account.

Once you've discovered that you can move your end mirror through a whole cycle of motion on a scale of a few μm, you might care to investigate how much *smaller* a motion you can detect. Your quadrature-detection scheme exhibits the greatest sensitivity to motion when you use the higher contrast of its two channels of output, and when you operate with the interferometer set at either of the two maximum-slope points: see the figure below --

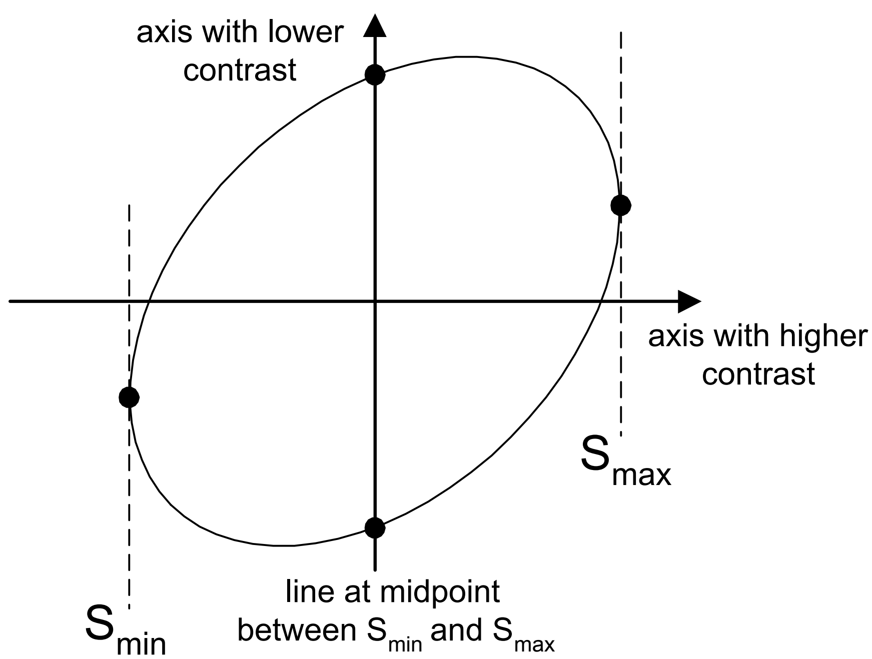

Figure 14-2: Operating points of maximum sensitivity in the XY-locus of quadrature Interferometry

You'll note that your signal changes from maximum to minimum for an actual mirror motion of $\lambda/4$, and is a sinusoidal function of mirror position between these extrema, so you can create a mathematical model for output signal as a function of mirror position x. From it, you can compute the sensitivity $\partial V_{out}/\partial x$ you can achieve by operating at one of the two maximum-slope locations.

Now you might try the direct depiction of the signal V_{out} on an oscilloscope while you drive the piezo stack with a time-dependent waveform, perhaps starting with an easily recognized triangle wave. You might want to use a wave with an offset, ramping up and down between extrema of (say) +13 V and + 1 V. If the amplitude is small enough, you'll see a faithful representation of the triangle-wave drive in the output of the interferometer. If the frequency is low enough (low compared to the mechanical resonances in the mass-and-spring system you are driving), the response will be frequency-independent in size; if you raise the driving frequency, the higher Fourier components of the triangle wave may start to excite mechanical resonances, with consequences you should be able to see and understand (and possibly also *hear*).

You'll note you might need by-hand intervention to keep the interferometer's state near the maximum-sensitivity condition, and you can provide these interventions either at your micrometer, or by using a dc offset of the waveform going to your piezo. After you've gotten good at this, you might try ever-smaller amplitudes of excitation to see how small an amplitude of motion of your mirror is still detectable above your noise level. A dual-trace oscilloscope display, with the 'scope triggered by the drive generator, with one channel displaying the drive, and with the second channel displaying the response, is the best eyeball way of looking for a weak signal. You might want to set your time base so that about 5 or 10 cycles of drive (and response) fit into your view. Your eye-brain pattern-recognition skills are quite adept at seeing the triangle-wave response, of known shape, period, and phase, even when sizeable amounts of noise are also present. If your 'scope has trace averaging capability, this is a great context in which to exploit its advantages.

This technique ought to convince you that mirror motions which are much smaller than $\lambda/4 \approx 0.16$ μm = 160 nm are directly detectable in real time. But if you know about lock-in detection, you can try for higher sensitivity still. You might change to sinusoidal excitation of the piezo translator, using the same generator's synchronizing output to provide the reference input to your lock-in amplifier. You might start with a sinusoidal drive of amplitude sufficient that you can see and quantify the response on a 'scope. When you capture that sinusoid on a lock-in, you will be able to understand the lock-in's response quantitatively (and also adjust its coupling, phase, and filtering settings properly). But now you can start decreasing the drive amplitude, perhaps by successive orders of magnitude, seeing if you can raise the lock-in amplifier's gain to restore your output. You will find that the lock-in will continue to display the steady presence of the signal even when your 'scope view shows the signal disappear down into the noise; you will find, as you go father along this path, that you'll need to use longer averaging times on your lock-in. But you should have no trouble detecting a regular sinusoidal motion of your end mirror of amplitude only 0.001 μm = 1 nm, and smaller still with the use of longer averaging times. If you can detect the motion of the end mirror with an amplitude of 0.1 nm, you will have the novel experience of 'seeing' the entire macroscopic assembly of stage and end mirror moving cyclically through an excursion of a size comparable to the diameter of a single atom!

If you have access to crystalline samples, you might see if you can detect the expansion of a bulk quartz crystal by this technique. Any rock shop or new-age store will offer natural quartz crystals, with convenient parallel faces already in place. You'll want to find a way to capture the crystal's thickness, perhaps between your micrometer's front face and the pushrod it bears upon; you'll also need to put some electrodes in place (one of them electrically isolated) to subject the crystal to a potential difference. You can be fully quantitative about the deformations you discover, and even without cutting any oriented slabs out of the quartz you'll have at least three distinct pairs of parallel faces to investigate. If you're worried about a 'control group' against which you can compare your quartz sample, you can substitute a piece of glass for the quartz sample.

15 The electro-optic effect

The electro-optic effect refers to a change in the optical properties of a material in response to an electric field. The simplest electro-optic effect to visualize is to imagine a slab of glass placed between a pair of capacitor plates, so that the glass can be filled with a static electric field; now you are to imagine a light beam passed through that slab, propagating in a direction perpendicular to the electric field. You should have some intuition that the optical properties of the glass might be affected by its electric polarization due to the static electric field; in fact, the index of refraction n of the glass will be changed into two indices, $n_{parallel}$ and $n_{perpendicular}$, applicable for light beams with propagation direction perpendicular to the static E field, but with optical E-fields that are either parallel or perpendicular to the static E field.

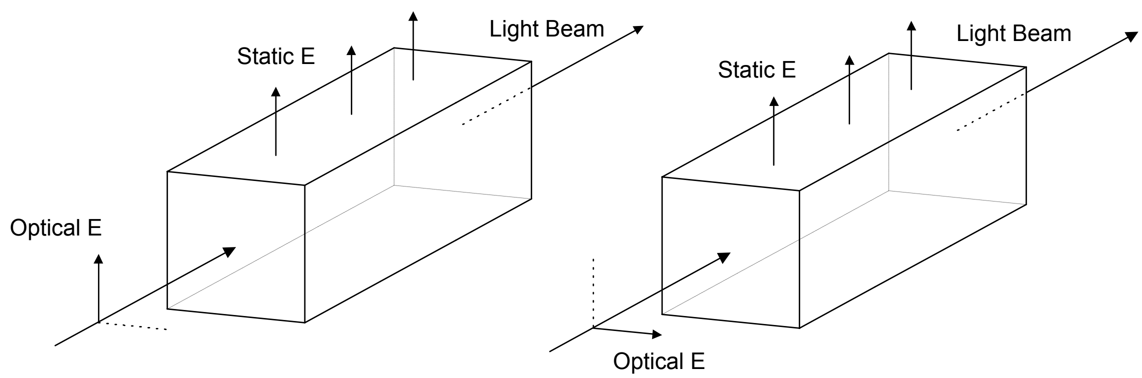

Figure 15-1: Two possible geometries for the static and optical electric fields inside a solid sample

a. The Kerr and Pockels effects

In isotropic media like glass, electro-optic effects are of second order in the electric field strength, which makes them hard to detect. [Look up the 'Kerr effect', and note that materials displaying the largest response are both carcinogenic and explosive!] So here we'll introduce you to *linear* electro-optic or Pockels effect, of the sort that can only occur in media *lacking* inversion symmetry. Your TeachSpin apparatus includes a sample of single-crystal lithium niobate ($LiNbO_3$), of size 8 x 8 x 15 mm^3, whose square 8 x 8 mm^2 end faces are optically polished, and whose top and bottom faces can be held between just the sort of capacitor-plate electrodes we mentioned above. Use the plastic tweezers provided for holding the crystal, touching only its roughened side faces, and store the crystal in its special pillbox when it's not in use. [Note that one of the sides of the crystal bears a black-dot symbol; for starters, let this face be a *side* face of your crystal as it's held in the holder.]

For detecting the small changes in refractive index that arise, the Sagnac interferometer is well suited. The geometry easiest to understand is the one with two separated beams; you should recall that these two beams propagate in opposite directions, and with orthogonal polarizations. The idea is to put the $LiNbO_3$ sample into one of these beams, while the other beam serves as a 'reference' or control-group beam. So confirm that you can get this sort of Sagnac interferometer to work, and put the two-plate tuner into place inside the interferometer so that you can smoothly scan through fringes. You'll want to use the polarimetric detection scheme so well suited to displaying the fringe signal that results.

Now put your electro-optic sample in its holder on an optical post-mount, and position it so that one of your two beams passes through the crystal. You'll be able to use the post-mount to adjust the vertical and horizontal positions of the crystal, to center it on one of your two beams. Confirm that the other beam still has free passage through the interferometer. You'll want the light to encounter the crystal's end faces perpendicularly; find a retroreflected beam arising from the end faces and use it, and the freedom of movement of the post-mount, to make this adjustment accurately.

After you've installed the crystal, you'll need to touch up the alignment of your interferometer (since the crystal's end faces might not be perfectly parallel). To do this, recall that you need to put a viewing screen *down*stream from the 45°-tilted polarizing beamsplitter cube in the polarimetric detector. Adjust the fringe pattern you see for zero spatial frequency, as usual, and then revert to ordinary electronic detection.

Now run the signal through a series of fringes, using the rotator mount on the two-plate tuner to move through a series of maxima and minima. What you're seeing is the result of differential phase shifts $\Delta\phi = \phi_{cw} - \phi_{ccw}$ accumulated by the beams propagating around the interferometer in the clockwise and counter-clockwise directions. If you record the values of the maxima and the minima, you'll be able to write down a mathematical model of your electronic output signal, of the form

$$S(\Delta\phi) = S_{max} \cos(\Delta\phi - \text{constant}) \quad ,$$

the point of which is that you can compute the sensitivity $\partial S/\partial(\Delta\phi)$ of your signal to small phase changes. This sensitivity is of course highest if you operate at points of maximum slope of your signal, which should be at or very near the zero-crossings of the signal. This is electronically very convenient, since you can raise the gain of your downstream electronics to as high a level as the noise permits. So now dial your two-plate tuner until you are at a zero crossing, and make an estimate of the noise or instability of your signal, and compute the size of optical phase shift to which this noise level is equivalent.

You may want to pay some attention to the draft shield, and to other ways to keep this noise level low (and hence your phase sensitivity high). When you're content with your interferometer's performance, you're ready to see if there are new phase shifts, and hence new signal changes, due to the electro-optic effect that is occurring in only one of your interferometer's beams.

A useful technique is to apply a square-wave time-varying potential difference to your capacitor plates, perhaps varying between -10 and +10 V, perhaps at a frequency of order 10 Hz or 100 Hz. If the same generator is triggering a dual-trace oscilloscope, and if the crystal drive voltage is displayed on one of the 'scope traces, you'll have the other trace free to display the Sagnac's output signal, which is directly sensitive to optical phase shifts. You're looking for that signal to display a waveform at the same frequency as the drive waveform, and in phase with it; your eyeball's ability to capture the presence of a small signal against a background of some noise is probably highest if the 'scope is displaying 5-10 cycles' worth of the waveforms.

Once you see any signal at all, you can confirm that your detection electronics have been set up with the right bandwidth to accommodate the rather fast changes you are making in your interferometer. You can also try varying your choice of drive frequency, to find the conditions under which the signal-to-noise ratio is optimal. [You're seeking a frequency at which the crystal response is still unaffected, but at which the noise signal is smaller.] You may also take advantage of signal averaging, if your 'scope offers that capability. Naturally, since you are looking for an ac signal of known frequency and phase, the technique of lock-in detection could offer you vastly greater sensitivity still, but here the goal is to achieve real-time detection on an oscilloscope.

If you can change your waveform from (0 to +10 V) to (0 to -10 V), you'll be able to confirm that this is the *linear* electro-optic effect: its sign depends on the sign, not just the magnitude, of the electric field. If you pay attention to whether you are operating on the ascending slope or the descending slope of the zero-crossings your two-plate tuner gives you, you can even assign the correct sign to $\Delta\phi$, the phase shift contributed by your E-O crystal. Finally, you can use the general relation

$$\phi = \frac{2\pi\, n}{\lambda_{vac}} L$$

for phase shift contributed by a sample of index n and length L to deduce from $\Delta\phi$ the size of the (truly tiny!) index change Δn that you are getting in response to the potential difference you are applying.

Now it's time for you to use a Polaroid to find out whether the laser beam that's passing through your crystal has its optical E-field parallel to, or perpendicular to, the vertical E-field your electrodes have created in your crystal. That's in preparation for sliding the crystal holder sideways, to get the light in the Sagnac's *other* beam to pass through it; that'll enable you to test for the other state of polarization. You may have to trim up the interferometer again, since your beam may be encountering a slightly different part of the crystal, but once your two-plate tuner can show you fringes again, you're ready to use the same electronic technique as before for seeing the $\Delta\phi$ signals, and deducing the index change Δn that result. You should be able to establish reliably that Δn is of the *other sign* than previously; with care, you should be able to assign the correct sign to each Δn.

With at least the magnitudes of the two index changes (Δn for parallel and perpendicular geometries) in hand, you can now arrange to see the net effect of both mechanisms occurring *simultaneously*. To do this, rearrange your Sagnac interferometer in the overlapped-beams mode, first taking out the crystal and realigning the interferometer, and repositioning the detector assembly, as required. You'll also need to substitute your quartz-plate tuner for the two-plate tuner in the rotational mount, to provide a way to dial through fringes in your overlapped-beams interferometer. When you have this working, re-introduce the E-O crystal into the (now overlapped) beams. You may want the crystal located about halfway around the interferometer from the input/output beamsplitter cube, since that's where the two counter-propagating beams are of about the same size. You may also want to work on the alignment, so the two beams are accurately overlapped in space there -- that way they'll both pass through the same area of the crystal, and they'll also be sampling the same air path all around the interferometer.

When you have the crystal in place, and the interferometer working at one of its zero-crossings again, you should cycle through some fringes to record the signal's maxima and minima again, to check again what sensitivity you have to small optical phase changes. You should also look at the noise level in your Sagnac signal, and ought to find it's quite a bit lower than previously: most of the noise due to uncontrolled index-of refraction variations in the air inside the interferometer is now common to both beams. Finally, you can look for the E-O signal, and now you hope to see good news in two forms: the signal level might be nearly double what it's been (if it's true that the index changes due to the E-O effect for the two polarizations of light are of opposite sign), and the noise level ought to be lower than it was (due to the 'common-mode rejection' you're getting with the use of overlapped beams).

With the improved signal-to-noise performance you're now getting, it would be well worth exploring how the net optical phase shift you're getting depends on the potential difference across the crystal. You do indeed expect to be seeing the linear E-O effect, but the slope of this line is also relevant. A figure-of-merit of the E-O device is its 'half-wave voltage', which is the potential difference that would need to be applied to get a net phase shift of π radians, or 180°, or half a cycle in phase. The TeachSpin crystal holder, and its cable and connector, are **not** rated for performance at this high a voltage, but there are applications that do use the half-wave voltage to execute optical switching or modulation. For example, if you've worked on the Mach-Zehnder interferometer, you know that a 180° phase shift suddenly introduced into one internal beam can cause the output light to flip from one output beam to the other. There are other topologies that can modulate a light beam's intensity, if not its direction, through the use of this E-O effect, and all such applications potentially profit from the extreme speed of operation possible in the E-O effect. The crystal's optical properties, and hence its effect on light, will change as quickly as you can change the voltages around it; typically the limiting factor is the rate at which some electronic driver can change these potentials, in the face of the capacitance of the plates-plus-crystal structure.

b. A 'collinear interferometer'

Once you've come to understand the E-O effect by using it on your separated or overlapped beams in a Sagnac interferometer, there is another topology in which you can use the E-O crystal as a modulator. It involves, in a sense, the use of two beams too, only this time they are overlapped in space and coincide in direction. Because the two beams are never separated at all, this new topology is maximally insensitive to air- and vibration-induced fluctuations. The two 'beams' to which I'm referring are the vertical, and the horizontal, polarization components of one single beam of light, linearly polarized at an angle of 45° with respect the vertical. I'll now show what effect the E-O device ought to have on that beam, and how polarimetric detection can best be used to detect that effect.

Suppose you set your HeNe laser source in its cradle mount with the 45° rotation you've used before, and use the usual two steering mirrors to produce a beam moving left to right across your breadboard. After the second steering mirror, you can install a Polaroid, with its pass direction also rotated at 45° with respect to the vertical, to ensure that you have a purely linearly polarized beam with a tilted polarization axis. Taking the z-axis as lying vertical in the lab, you could write

$$\mathbf{E}_{in} = E\,\left(\frac{\hat{\mathbf{y}} + \hat{\mathbf{z}}}{\sqrt{2}}\right) \exp i(k\,x - \omega\,t) \;,$$

where the point is that the y- and z-components not only have equal amplitudes, they now are sure to have the same phase too. Now send that beam directly through your E-O crystal in its mount (no interferometer at all!) and note that the output state of the light can be written as

$$\mathbf{E}_{out} = \frac{E}{\sqrt{2}} (\hat{\mathbf{y}}\ e^{i\,\phi_y} + \hat{\mathbf{z}}\ e^{i\,\phi_z}) \exp i(k\,x - \omega\,t) \quad ,$$

where ϕ_y and ϕ_z describe the (different) phase shifts that the crystal might have caused for the two distinct polarization components in the beam:

$$\phi_y = \frac{2\pi\, n_{perp}}{\lambda} L \ ;\ \phi_z = \frac{2\pi\, n_{para}}{\lambda} L \quad ,$$

since the static electric field in the crystal is in the z-direction. Insofar as these two phase shifts are different, the emerging light is no longer linearly, but rather elliptically, polarized. The question is -- what is a sensitive and efficient way to detect this? The answer is to use the same polarimetric detection scheme as at the output of the Sagnac interferometer, which uses a polarizing beamsplitter cube oriented at 45°.

By projecting the field $\mathbf{E}_{out}$ onto the two unit vectors

$$\frac{\hat{\mathbf{y}} \pm \hat{\mathbf{z}}}{\sqrt{2}}$$

which describe the polarization eigenstates of the PBSC, and then computing the beam intensities that will be sent on to the two detectors, we obtain

$$S_+ = \frac{E^2}{4} (1 + \cos \Delta\phi) \ ,\ S_- = \frac{E^2}{4} (1 - \cos \Delta\phi) \ ,$$

where $\Delta\phi = \phi_y - \phi_z$ is the phase difference between the two components of $\mathbf{E}_{out}$. Now the difference between those two signals, easily obtained electronically from the controller box, has the desirable feature of being centered around zero, with form

$$\Delta S = S_+ - S_- = \frac{E^2}{2} \cos \Delta\phi \quad .$$

But (because of the cosine dependence) this does *not* provide a signal that is linearly sensitive in the presumably small phase shift $\Delta\phi$ you're trying to detect. In fact the problem is more complicated than that, because even in the limit of zero electric field on the E-O crystal, the phase difference $\Delta\phi$ might not be exactly zero. The cure for that second-order problem, and this birefringence problem too, is to add the quartz-plate tuner unit somewhere between the linear polarizer and the PBSC. There it will create its own, variable, and conveniently adjustable, differential phase shift $\phi_y - \phi_z$ between the y- and z-polarized components of the beam. Then the *whole* phase difference, due to the crystal's birefringence, due its additional electro-optic effect, and due to the quartz plate, can be written as

$$\phi_y - \phi_z = (\phi_y - \phi_z)_{crystal} + (\phi_y - \phi_z)_{E\text{-}O\ effect} + (\phi_y - \phi_z)_{quartz} \,.$$

The empirical method is to use the rotation stage to adjust the quartz plate's contribution to cancel most of the crystal's contribution, leaving a difference which is just $\pi/2$ (modulo π; then the overall optical phase shift difference becomes

$$\phi_y - \phi_z = \pi/2 + (\phi_y - \phi_z)_{E\text{-}O\ effect},$$

and the signal emerging from the polarimeter becomes

$$\Delta S = \frac{E^2}{2}\cos\Delta\phi = \frac{E^2}{2}\cos\left(\frac{\pi}{2} + \Delta\phi_{E\text{-}O}\right) = -\frac{E^2}{2}\sin\Delta\phi_{E\text{-}O}\ .$$

Operationally speaking, we're just using the quartz plate to bring the observable signal to one of its zero-crossing points again, and there we have the usual payoff of getting *first*-order sensitivity to the small electro-optical phase shift that the crystal can create. The usual rotation of the quartz plate also affords the chance to change the sine-function's value to ±1, thereby giving an empirical value for its coefficient, and a method for establishing the sensitivity of the method.

So how does this method differ from the use of the Sagnac interferometer? Optically speaking, not very much; in both arrangements, you can use the same Polaroid for preparing a linear polarization, the same E-O crystal and quartz plate for modifying that polarization, and the same polarimetric detector for analyzing the final polarization. In the Sagnac interferometer, you have the possibility of separating the two polarization components into spatially separate beams, which is crucial to the separate measurement of $\Delta\phi_y$ and $\Delta\phi_z$. You also have the incidental difference that the two polarization components are moving in counter-propagating directions. In the non-Sagnac case, you have the two polarization components present in one single beam, guaranteed to follow the same path through the optical components and the air. So the Sagnac method might be conceptually simplest, but the non-Sagnac variation might give signals of the lowest possible noise.

c. Conoscopic figures

Here's a wonderful demonstration of a classical optical property that you can view with your lithium niobate crystal; it has nothing to do with the electro-optic effect, and little to do with interferometry, but it gives much too pretty a picture to pass up.

The idea is to view the crystal through crossed polarizers, and you have all the tools you need to get a great view. For the best visual impact, you'll want to use your white-light source, and a sheet of white paper serving as a diffuser screen to give a featureless white light-emitting source. Now arrange for light leaving the diffuser to pass through a 'sandwich' of a linear polarizer, the crystal, and a second polarizer. Set up your sandwich with the second polarizer, and the crystal, as close to your eye as you can achieve; this gives you the widest possible field of view into the crystal.

Look toward the white-light source and enjoy the view. See how it changes as you rotate the polarizer(s), and see if you can identify the center of a pattern. (You may want to rotate the crystal mount on its base to optimize that view.) By viewing white light that passes through both polarizers but bypasses the crystal, you have an easy way to determine the state of the polarizers. Now look a bit more quantitatively--

- when the polarizers are 'crossed', what's the state of the center of that pattern?
- when the polarizers are 'aligned', what's the state of the center of that pattern?

The 'center of the pattern' defines a direction within the crystal, not a position inside it -- how can you tell that? This direction in the crystal is called the 'optic axis', and like the concept 'north', it is a direction and not a position. See if you can arrange the tests that convince you that light of any given angle of linear polarization, propagating through the crystal in the direction of the optic axis, emerges from the crystal with its linear polarization unchanged.

Of course the beauty of your view arises from seeing light that comes through the crystal along directions *other* than that of the optic axis. You can see the circularity of certain features in your view; such circles represent whole cones of rays inside the crystal, all making the same angle with respect to the optic axis. In fact your view is called a 'conoscopic figure' because you're seeing the effects along such cones.

What's going on here? The explanation of all that you see is based on the *birefringence* of lithium niobate. For light propagating along the optic axis, with an E-field perpendicular to the optic axis, the material is characterized by an ordinary refractive index, n_o, which is the *same* for (say) vertically as for horizontally arranged E-field. But for light propagating at some internal angle θ with respect to the optic axis, there are *two* indices of refraction: the ordinary (and θ-independent) n_o, and a new 'extraordinary index of refraction' labeled n_e (which does depend on angle θ).

You can read about the theory of n_o and $n_e(\theta)$ in any treatment of birefringence, and you could investigate both indices quantitatively by interferometry, using a monochromatic beam passing through a crystal which you can rotate by a controlled amount.

But what creates your wonderfully colorful 'conoscopic figure' is the *wavelength* dependence of both indices of refraction. Perhaps you can intuit, or devise a way to view, a conoscopic figure for monochromatic light. If so, perhaps you can intuit how this monochrome view would change as the wavelength λ changed. (Hint: on axis, there's no λ-dependence; but off axis, circular features would grow or shrink as λ changed).

Now you can put together a picture of what creates the colors in your conoscopic figures. Suppose, with polarizers crossed, you're looking at light on the cone at angle θ where (say) green light propagates exciting both the 'ordinary' and the 'extraordinary' refractive properties of the crystal. With two distinct indices $n_o(\lambda_{green})$ and $n_e(\lambda_{green}$, this $\theta)$, you can imagine that the phase accumulation $\phi = 2\pi\ n\ (L/\lambda)$ is different for the o- and the e-waves. The result is that green light can emerge from the crystal with a *different* polarization than that with which it entered, and thereby some of it can get through the second, crossed, polarizer.

With this picture, it's fun to give an account of the colors that appear in this conoscopic figure. We'll do so in the case where the polarizers are crossed, so the pattern's center is black. We'll use the language of the 3-receptor model of human vision, and pretend that the white light incident on the crystal is made of just red, green, and blue (RGB) monochromatic components.

You see that on axis, the crystal acts like an isotropic material, and all three colors (R, G, and B) have their polarization unchanged by the crystal and so are blocked by the second polarizer. The visual effect is <u>black</u>. Now going off-axis in direction, the mechanism of birefringence produces light emerging from the crystal for which R, G, and B are all changed in polarization, such that all three can pass through the second polarizer. The visual effect is <u>white</u>.

Farther off axis, the polarization effects continue to increase, and because of the crystal dispersion (higher indices of refraction toward the blue end of the spectrum), the blue light is the first to return to the state of blockage that it had at the pattern's center. For this angle, with B blocked, and some transmission of R and G, the visual effect is yellow.

Farther still in angle, some blue light reappears, but by this angle, it's green light that is blocked. With G blocked, and some transmission of R and B, the visual effect is magenta.

Yet farther in angle, we get to the first occasion (apart from the central axis) for which red light is blocked. With R blocked, and B and G transmitted, the visual effect is cyan.

The color sequence becomes more attenuated in hue, and complicated in sequence, for larger angles off axis. You can look for the first place where red and blue are both (nearly) blocked, with the visual result of nearly pure green. Or you could try to send light from a given part of the pattern into a spectrometer, to see the way in which polarization blockages for certain colors appear as minima in the spectrum. Whether you analyze this effect visually, qualitatively, or quantitatively, you're sure to reach a better understanding of the RGB-model of human vision, and the theory of complementary colors.

There are observable features not yet addressed by the discussion above: if it describes the circular features whose colors you've been noting, what accounts for the Maltese-cross pattern overlaying them? Furthermore, this model of a 'uni-axial crystal' assumes that the crystal's optical properties are cylindrically symmetric around the optic axis; what about the symmetry of its electro-optic properties? Try rotating the crystal about its long axis by successive 90° rotations, and see what consequences you can find in the optical, and electro-optical, properties it displays.

16 White-light interferometry

Perhaps you've read some textbook introductions to lasers that emphasize that lasers produce 'coherent light', unlike all other sources, which produce 'incoherent light'. Perhaps you've gotten the impression that experiments in optics (like interferometry) depend on the use of 'coherent light'. The central lesson of this section is that there is only one kind of light, and coherence is a matter of *degree*, not *kind*. In particular, interferometry provides a fine way to investigate the *degree* of coherence of light, and white-light interferometry shows that even an incandescent bulb's light output possesses a sufficient degree of coherence to display interference phenomena like fringes.

The apparatus best suited for the initial investigation of coherence phenomena is a two-arm Michelson interferometer, using the 0-1" translation stage to support one end mirror, and the flexure stage to support the other end mirror. The metal-film beamsplitter, as in the section on quadrature interferometry, is particularly useful for tracking down the central requirement of white-light interferometry, which is to match the optical lengths of the two arms of the interferometer not just to the nearest inch or millimeter, but to a tolerance of about 1 μm. Naturally, this matching will be achieved using interferometry itself as a testing tool.

One more special requirement to white-light interferometry is to ensure that the equal-arm condition is simultaneously achieved for a whole range of wavelengths. This would be easy enough to achieve if beamsplitters acted the way they are sometimes drawn, namely as infinitely thin geometrical planes. But real beamsplitters are films deposited onto substrates of glass, and the drawing of Fig. 16-1 below (shown, for clarity's sake, with a misaligned interferometer) shows that a thickness of glass substrate which is traversed only *once* by one beam is traversed *three* times by the other beam.

You'll note that, for this drawing to be useful, it's necessary to know which side of the glass substrate of the beamsplitter in fact bears the metal-film coating. This is best confirmed by removing the beamsplitter optic from its mount, and viewing its roughened edge; there's a penciled feature at one point on the periphery, looking like < or >, standing for an arrow pointing to the face bearing the metal film. You will want to read Appendix H, on 1" optical mounts, for details on this procedure; whatever else you do, be sure you handle the 2-mm thick optical slabs by their *edges only*.

Since the optical thickness of glass depends on its refractive index, and since that refractive index varies with wavelength, the desired arm-length matching condition can*not* be met over a whole range of wavelengths. [The idea is that an extra (say) 4 mm of glass in one arm could be matched -- at any one wavelength -- by a required number of mm of air in the other arm; but that this matching won't work simultaneously over a range of wavelengths, since the dispersion of glass mismatches the dispersion of air.] The cure is a 'compensator plate', made out of the same thickness of the same kind of glass as the beamsplitter, and installed (also at a 45° angle) in the beam which otherwise passes through only one thickness of glass. [Of course the compensator plate is anti-reflection coated on *both* faces, as the beamsplitter is AR-coated on its *second* face.]

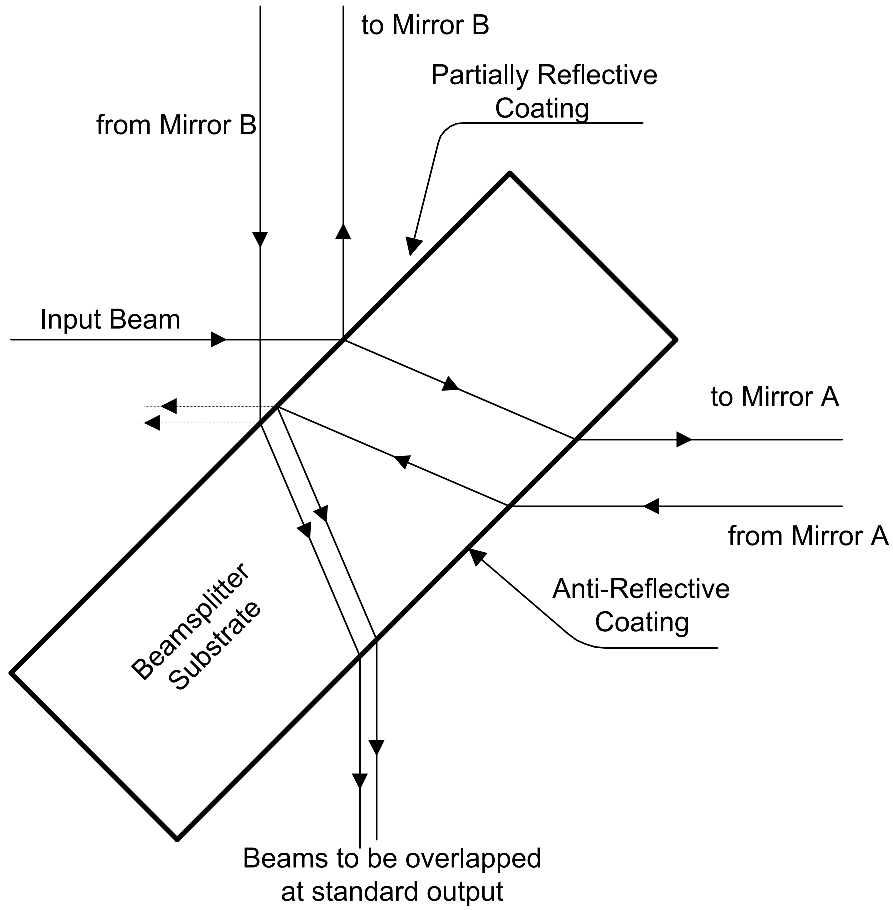

Figure 16-1: Light paths in a Michelson interferometer's beamsplitter, **not** ignoring the thickness of the substrate of the beam splitting film (paths shown separated for clarity)

In practice, it's far from trivial to ensure that the thicknesses of the beamsplitter and the compensator are matched to 1 μm; happily, the effective thickness of the compensator can be fine-tuned by rotating it a trifle from the 45° condition. Instrumentally, you can achieve this by mounting the compensator plate in a 1" lens-holder on a post in the rotational stage you've used in Section 11b. For starters, you can use an eyeball-accurate setting to the 45° condition, and set the rotational stage to accommodate a future adjustment either way by a few degrees.

Finally, the adjustment of your interferometer to the equal-arm condition will be easiest if you install a second (50-50 dielectric) beamsplitter, as in quadrature interferometry, upstream of the interferometer in the input beam so as to monitor the 'non-standard output' of the Michelson interferometer.

Now that you have all the parts in place, you will want first of all to align your interferometer using the HeNe laser as light source. You should look not only for two spots to overlap and interfere at the standard output plane, but also for the non-standard output to be heading nearly straight back toward the HeNe laser. It might take a little shimming of the mirror mounts to achieve this. Once you have fringes, you can adjust them for zero spatial frequency as usual, and confirm that adjusting the micrometers on either end mirror gives you a smooth alternation of bright and dark in the output spots as you expect. You'll want to use the differential micrometer, together with a pushrod, with the flexure mount, just as in Experiment 9, in order to get resolution and control, at the 1 μm level, over the path difference in the interferometer.

Now the interferometer you've just set up will give good fringes even if the path difference in the two arms is a millimeter, or a whole inch; here's the procedure for 'zeroing in' on a path difference which really is zero to a 1 μm tolerance.

1) You'll want first to use the witness mark atop the beamsplitter holder as a reference point from which to measure the distance to each end mirror, to a 1 mm tolerance. You can ignore the offset, and the refraction, in both the beamsplitter slab and the compensator slab [provided, of course, that you have installed the compensator in the correct arm of the interferometer!]. Once you've measured both arm lengths, you can use the 0-1" translation stage to equalize the arm lengths to better than the nearest mm. You should at this stage have the flexure mount centered in its ±1 mm range of motion, and the diff mike similarly centered, so that you have available the displacement you'll need to achieve the exact arm-length matching condition.

2) What you need now is a *diagnostic* for arm-length inequality, and this is best provided using a monochromatic source of adjustable wavelength. A red-wavelength tunable dye laser would work perfectly, but it's a very expensive accessory, so the TeachSpin kit provides a much cheaper surrogate, which is a thermally tunable red diode laser. You're finally going to have a use for the two thermoelectric modules (TEMs) mounted on the circumference of the metal cylinder holding your diode-laser module. You must ensure that the two TEMs have their active faces in contact with the 45° tilted faces of the aluminum upright of the source mount, since the aluminum upright will serve as the heat sinks required for the operation of these solid-state heat pumps. You should also be aware that the TEMs' faces are ceramic, and about as brittle as glass, so you need to be careful in handling this diode-laser source, and careful not to over-tighten the clamping thumbnuts in the source holder. Once you have the diode-laser source mounted, you can connect it to the Modern Interferometry controller box via its cable, and activate it using the front panel switch. You should see the red beam, nominally 650 nm wavelength, emerging. What's novel is that the back-panel accessory inputs now provide a way to send dc electrical current to the TEMs, and the back panel toggle switch provides a way to reverse the direction of this current at will. This allows you to use the TEMs alternately as thermoelectric coolers and heaters.

3) You'll need an external adjustable dc power supply, capable of providing up to 0.5 A, to connect to the back-panel connectors to power the TEMs. You'll find it takes a few Volts to drive 0.3 or 0.5 A through the TEMs. You'll see very little effect of this TEM current, unless you also monitor the temperature of the diode laser, or its output wavelength. It might be easiest to send the diode-laser beam through a transmission grating, and follow the diffracted beam to a screen a meter away, to observe the effect on wavelength. Upon first turning on the TEMs, you might see the diffracted spot start to move. You can expect the output wavelength of the diode laser to increase by approximately 0.2 nm per K of temperature increase; you should see the wavelength equilibrate to a new value with a

timescale of order 30 seconds. Then if you reverse the direction of the TEM current, you should see the wavelength shift in the other direction.

[If you use a higher-resolution spectrometer, you'll see more details, including the multi-mode structure of the diode laser, and the 'mode hopping' that occurs as you change the temperature. Typically, you'll see individual modes move in wavelength at approximate rate +0.05 nm/K, and the median wavelength increases, discontinuously, at the approximate rate +0.2 nm/K. You'll see that the diode laser is *not* monochromatic, nor continuously tunable, but the important thing is that you <u>can</u> vary its wavelength, or mix of wavelength composition, on demand, by reversing that TEM switch.]

4) Why does varying the diode-laser wavelength matter? Go back to the case of of an unequal-arm Michelson interferometer, with (one-way) arm lengths of L_1 and L_2, and understand the derivation that predicts an output signal of

$$S(f, \Delta L) = S(f, 0) \cos^2 \frac{2\pi f}{c} (L_1 - L_2) \ ,$$

where f is the frequency of the input light. Ordinary operation of such an interferometer is to fix f and vary (say) L_1, in which case you expect to get fringes. Here you're fixing L_1 and L_2, and instead varying the frequency f. You will get fringes, <u>except</u> in the special case of $L_1 = L_2$, the equal-arm condition you're trying to reach. You might now work out what frequency change you get for a (say) +0.5 nm change in wavelength near 650 nm, and what fringe signal this gives for a path difference of (say) $L_1 - L_2 = 1$ mm.

5) The best diagnostic is to set up two photodetectors and a 'scope in the XY-mode to view the elliptical locus of the two signals (from the standard and non-standard outputs) of your quadrature Michelson interferometer. You can fine-adjust the alignment of the interferometer to maximize the contrast of the two signals, and then let the locus stabilize to a point (exhibiting the usual sensitivity to vibration, and now a new sensitivity to the drifting wavelength of the diode laser). You'll need to install the draft shield to get the stability you'll require for the next tests. The goal is to activate a temperature excursion of the diode laser, up and down in a slow cycle, by reversing the TEM currents perhaps every 30 s, and to note how the point moves along the locus in response to the resulting wavelength changes. The prime convenience of quadrature detection is that you'll be able to say *which way* along the locus the point moves (clockwise or counter-clockwise) for a temperature increase.

Now you might use the differential micrometer to increase one arm length by 0.25 mm, and repeat.

- Has the effect (of diode-laser temperature changes) on the fringe signal changed direction? [If so, you have changed the *sign* of $L_1 - L_2$, which means that you have gone past the equal-arm-length condition.]
- Has the effect on the fringe signal increased in rate? [If so, you have increased the magnitude of $L_1 - L_2$, and should make the next adjustment in the other direction.]
- Does the fringe signal collapse at times, so the elliptical locus smears down to some fuzzy blurs? [If so, you are certainly not at $L_1 = L_2$, and you're also seeing the diode-laser pass through a mode-hop condition during which it's emphatically not monochromatic.]

Your goal is to use this diagnostic to get to a point where temperature, and hence wavelength, changes in your diode-laser make no discernable changes in the location of your signal point on your locus. You should be able to show that this requires the arm lengths to be matched to better than 0.1 mm or 100 μm.

6) You can now make two changes in your interferometer -- you're ready to switch to a light source whose wavelength you can vary much more; that's the red/green LED source unit. So dial down the TEMs' power supply to 0 Volts, remove the TEM-equipped diode-laser source, connect the red/green LED source instead, and mount it in the source holder in place of the diode-laser unit. Dial up your power supply to about 12 Volts, and you should see the light-emitting diode light up, either red or green. Activating the toggle switch will still reverse the direction of the current, and will operate the other-color LED in the single housing that bears them both. You might monitor the current being supplied, and set the power supply so that a current of ±20 mA flows through the LEDs. Note that this (red or green) LED light is originating from the same place that the HeNe or diode-laser beam formerly did, and note that some part of that light is propagating into your interferometer.

The other change you'll need is to install a Polaroid somewhere between the LED source and the interferometer's entrance, in order that the interferometer be illuminated with linearly, and vertically, polarized light. Your previous laser sources were set up to deliver properly polarized light already, but since your interferometer with metal-film beamsplitter delivers decent fringe contrast only for vertically polarized light, you'll need the Polaroid to ensure that in the case of the LED source.

How will you detect the output of your interferometer, now that there's no collimated output beam? **In these *non*-laser cases only**, you may put your *eye* into the place formerly taken by the photodetector for the non-standard output of the Michelson interferometer, and *look* right into the interferometer. You should see, as if at a distance, the LED-as-source in your field of view. But better to see the fringes you're looking for, it's important to build a little telescope by which to sight into the output beam of the interferometer. Curiously enough, you want to focus this telescope not 'at infinity', nor on the LED-as-source, but rather on the (common) optical plane occupied by the end mirrors of your interferometer. You can temporarily stand some paper-with-printing up against one end mirror and practice aiming and focusing with your telescope; of course, you'll need to understand which two plano-convex lenses to use to give you the magnification desired.

7) The claim is that for the proper setting of the differential micrometer, you will see your view of the mirror plane modulated by some pattern of constructive and destructive interference curves. Over a small range of diff mike settings, centered on the zero-path-difference condition you are searching for, your view of the mirrors should show fringes in space. To find these features, you'll have to adjust patiently the diff mike over a range of perhaps ±0.1 or ±0.2 mm, the 100 or 200 μm that may yet be separating you from the equal-arm condition. Once you see any fringes, you'll know they're real if a mere adjustment by 20 μm or so will cause the fringes to decrease in contrast down into invisibility, and if the fringes then come back again upon returning to the 'magic' setting of the differential

micrometer. Blocking the light in either arm of the interferometer should also obliterate the fringes.

8) Once you see these interferometric fringes, read the micrometer setting at which they attain optimal contrast, and then use the toggle switch to change the color of the LED's output. Use the diff mike to locate the same kind of fringes for the other color, and record its new position setting. If your compensator plate is doing its job, the locations should be similar, to perhaps 10 μm; you can now use these readings to adjust the compensator plate more perfectly. You might rotate the plate by 1°; then you'll have to use the diff mike to search, over perhaps 50 μm, to find the fringes again. Again, read the settings needed to get optimal-contrast red, and green, fringes. If those reading are closer together, you've improved the degree to which the compensator plate is working. Use this diagnostic until the red and green fringes occur at the *same* diff mike setting, to perhaps 5 μm tolerance.

9) When you have seen, by direct viewing, the fringes in red and green LED light, you are finally ready to make one last light-source substitution, by changing to the incandescent bulb source. This will again produce light at the same place as your laser and LED sources, and some of it will head through the interferometer and reach your eye. Dial the power supply up from zero, still driving your source via the accessory connections on the controller's back panel, until you get a comfortable level of light into your eye. You'll get the best eyeball sensitivity to colored features if you use a rather dim setting of the bulb. If you have adjusted the interferometer properly using your LED source, you should already be able to see curiously colored fringes in your telescope's view into the interferometer. You can look for the maximum-contrast condition of your interferometer by fine adjustments of your diff mike; with this light source, changes in one arm length of just a few μm will take you right off the high-contrast parts of your interference effect. You should now note that the fringe of highest contrast is nearly pure black-and-white, while adjoining fringes take on a series of colors. You should be able to explain all that you're seeing in qualitative and even quantitative detail; but whatever else you've accomplished, you have seen a proverbially *in*coherent source of light nevertheless giving fringe signals in a Michelson interferometer! You've also used an optical diagnostic which picks out one unique micron of translation of a given mirror from all the $>10^5$ microns or 100 mm over which you could imagine translating it.

10) Once you've found white-light fringes like this, you might measure how many microns of translation of one end mirror suffice to take you away from visible fringes. Once you've done that, you'll want to go back to the central fringe of the fringe pattern, and then watch the same fringes, only this time viewing them through a glass filter. Your equipment includes two small green 'interference filters', and you might first use the subjectively more transparent one, which passes wavelengths in a band 80 nm wide centered on 550 nm. You might first mount that filter *up*stream of your interferometer. Of course you'll see dimmer fringes, and now they'll be more nearly single-color in appearance; but try again to see by how many microns you can translate the end mirror before the fringes disappear for lack of contrast. I predict you'll be able to go *farther* from the 'central fringe' before the fringes disappear, and I further predict that the more narrow-band the filter, the farther-still from

central-fringe you can go and still see some fringe contrast. To test that claim, try using the other green filter, which passes wavelengths in a band only 10 nm wide, centered on 546 nm.

One more surprise: I claim that the effects of these filters on your view will be the same if the filters are mounted *down*stream of your interferometer, inside or even beyond your telescope. You'll profit from thinking hard about how a filter 'after the fact' can seem to change your interference pattern.

You've been seeing the direct trade-off between breadth-of-spectrum and coherence-length that is captured in the following derivation of the signal that's expected in an interferometer:

> if s(f) df gives the strength of the signal, at the detector, attributable to frequency components between f and f + df, in the *absence* of interference, and s'(f) df gives the strength of the signal, at the detector, attributable to frequency components between f and f + df, in the *presence* of interference, where
>
> $$s'(f)\,df = s(f)\,df \cdot \cos^2 \frac{2\pi\,\Delta L}{c} f \quad ,$$
>
> then, given a detector which responds to all frequencies, its total output would be
>
> $$S = \int_0^\infty s(f)\,df$$
>
> in the *absence* of the interferometer, and it will be
>
> $$S' = \int_0^\infty s'(f)\,df = \int_0^\infty s(f)\,df \cdot \cos^2 \frac{2\pi\,\Delta L}{c} f$$
>
> in the *presence* of the interferometer.
>
> You'll note that the signal S' = S' (ΔL) is certainly a function of the path-length difference $\Delta L = L_1 - L_2$; in fact, this signal, recorded systematically as a function of ΔL and plotted, is called the 'interferogram' of the source. You should be able to show that subtracting the average value from the interferogram, and then taking a Fourier transform, produces as output the spectral distribution s(f) of the source (perhaps as modified by the transmission-efficiency-versus-frequency of the interferometer optics, and the detection-efficiency-versus-frequency of the detector). You've just extracted the spectrum with*out* using anything like a prism or grating! This is the basis of 'Fourier transform spectroscopy', and it is not just an academic curiosity; rather, it is the actual technical basis of most of the spectroscopy that's nowadays performed in the infrared region of the electromagnetic spectrum.
>
> Your interferometer's source, solid-angle of acceptance, and detector are not optimized for practical Fourier-transform spectroscopy, but you should be able to deduce lots of qualitative results from the Fourier transform relation which connects the spectral distribution and the interferogram. In particular,
>
> - a source whose spectral distribution approximates a delta-function in frequency yields an endless sinusoid for interferogram;
> - a source whose spectral distribution approximates a flat function in frequency yields an approximate delta-function for interferogram;

- a source whose spectral distribution has a width Δf around frequency f_0 yields an interferogram with a width of order ΔL around zero path difference, where the product $\Delta f.\Delta L$ has a value of order c.

This trade-off ought to remind you of the so-called 'uncertainty principle' of wave mechanics, and it effectively defines the concept of 'coherence length' of an optical source. This is also why your incandescent source, which together with your interferometer optics and your eyeball-detector sensitivity define a spectral distribution covering from perhaps 400 THz to 600 THz, give even 'white light' a frequency width of $\Delta f \leq 200$ THz, and hence produce visible fringes over a path difference of order c/200 THz or about 1.5 μm. This trade-off also accounts for the larger range of translation ΔL that will give fringes when you restrict, by filtering, the frequency distribution giving rise to the detector signal.

11) Here's a modification you can make to your white light interferometry for the sake of a more visible display. Thus far, you've seen the features are modulations in the eyeball-received light signal, as a function of the setting of the path-length difference. To get a view of fringes-in-space rather than these fringes-in-time, you need a view in which an entire (small) range of path-length differences is present all at once. You can achieve this by mere tilting of one end mirror, if you can ensure that you don't also translate an end mirror by more than a few μm. So find the end-mirror whose thumbscrew adjustment bears on a flexure with 'horizontal hinge line', and notice that the pivot point of the mirror is very nearly directly behind the center of your light beam. So turning this thumbscrew will tilt the mirror, but the active part of it will scarcely translate at all. [By contrast, the mirror with 'vertical hinge line' will, upon adjustment of the thumbscrew, not only get tilted, but will also have its relevant part *translate* by many μm, taking you away from the equal-arm condition.] So watch the interferometer output by eye as you carefully make very small adjustments of the correct thumbscrew; you should see a series of fringes-in-space overlay your view of the source. You may need to make fine adjustments of the diff mike to keep the 'central fringe' centered in the illuminated part of your view. You're getting a view, by white-light interferometry, of the 'virtual wedge' formed by the two mirror surfaces, which 'virtually intersect' on a zero-path-difference line, and give you a view in space over a whole range of arm-length differences. Now you get an 'interferogram in space' instead of one spread out in time as you pass through a range of path-difference values.

Now you can try changes to either end-mirror's thumbscrew adjustment, and watch the effects on the white-light fringes. Ideally, the colored fringes would be straight and parallel, and would turn into a uniform white field of view when the wedge angle were set right to zero (and the diff mike were set to zero path difference). In practice, you're seeing the *failure* of your optical components to be accurately flat over your full field of view. Perhaps now you can see why interferometric methods can be used to test the difference in figure between a reference surface and an attempted duplicate of it. Perhaps you can already understand that your pattern's central (most color-free) fringe can serve as a 'contour line' of zero-path-difference between the surfaces of the interferometer's two end mirrors, and you can see how that would allow a 'test object' and a 'reference object' to be compared with high sensitivity.

12) A second and more ambitious change you might try to make is to change from the metal-film to the dielectric-layer beamsplitter in your interferometer. The idea is to *avoid* having to search again for the zero-path-difference condition, by changing *nothing* but the beamsplitter; this certainly means that you'd *not* unbolt the beamsplitter mount from the optical breadboard to do the swap. Instead, you'll have to unscrew the threaded conical insert holding the 1"-diameter beamsplitter plate and remove the one plate, and install the other, in its place. It's not trivial to do this, especially so since you want to avoid touching either beamsplitter plates' optical surfaces. You'll also have to ensure that the new beamsplitter gets put into its mount with its 50/50 dielectric layer placed on the same *side* of the glass substrate as was formerly occupied by the metal film. (See Appendix H on identifying the distinct faces of beamsplitter plates.)

If you can accomplish all this, the active side of the beamsplitter will nominally take up the same position in space as the one it's replacing; but in fact the flatness of the beamsplitters is imperfect, and you'll have to scan by perhaps ±10 μm to retrieve the position of zero path difference. Once you find it, you can remove the upstream 'other' beamsplitter, since you can now longer accomplish quadrature detection anyway. You can also remove the Polaroid in your illumination path, as your new interferometer is polarization-insensitive. The combination of changes will markedly increase the intensity of the light you are receiving, and may also improve the contrast of your fringe pattern. The goal is to get a showpiece central fringe that's really white, with adjacent interference minima that are as dark and black as possible. [You may need further rotational trimming of your compensator plate, and accompanying re-adjustments of your diff mike, to optimize this compensation.] Fringes farther from the central fringe will take on a delicate series of shaded colors, which you should be able to understand via the principle of superposition -- you're seeing the interferometer work simultaneously for all the colors present in the light source and transmitted through your optical train.

With this modification, the sensitivity of your interferometer is high enough that you can view fringes from rather dim or diffuse sources. A great example is to turn off your white-light source entirely, and place in front of it a plain white card. Now if you illuminate that card with any light source, it will scatter enough light into your interferometer that you can view the interferogram of that light source. A compact-fluorescent bulb yields a fascinating interferogram, and a view of its light through a spectrometer might tell you why you get such a surprising outcome.

13) There are still two more ways to display your fringe patterns. If you'd like to project them on a screen, note that the upstream lens of your telescope ought to be generating a real image of your 'localized fringes'; that image will be located in space somewhere between the two lenses in your telescope. So you might want to dim the room lights, and dial up your white-light source (staying under its 1.1 A current limit), and then move an alignment tower or other screen along the beam path inside your telescope. At a certain plane in space, you'll see your fringe pattern in focus, projected on the screen. You can adjust your 'mirror wedge' in such a way that light coming through the little hole in your alignment tower takes on a lovely succession of hues as you adjust arm-length differences.

If you'd like to *photograph* your fringe pattern, you can get rid of your telescope altogether, and just use the *camera's* lens (instead of your telescope's upstream lens) to focus the pattern to form a real image onto the film or CCD in your camera. You might want to use a moderate telephoto lens (since the pattern is of small angular scale), and you might want a telephoto lens with a 'macro' setting (since you have to focus the camera not at infinity, but rather on the relatively nearby plane occupied by the interferometer's end mirrors).

Appendix A. Suggested sequences for experimentation

There are so many investigations that can be accomplished using the tools in your Modern Interferometry kit that it's worth suggesting some logically connected *subsets*.

Any investigator ought to start reading Section 0, and ought to get familiar with the basic components and operation of the kit by performing the basic Michelson interferometry of Section 1. Parts of Section 2, on Optical Alignment, will be useful to read and experience also. With just these techniques in hand, the Experiments of Sections 9 and 11 are accessible.

Beyond that point, there are some natural subsets for investigation.

- Section 3, on the Quadrature Michelson interferometer, is interesting in its own right, and also provides the tools that make measurements of Sections 10, 12, 13, or 16 feasible too.

- Section 5, on the Sagnac interferometer, is an alternative. It will require, and reinforce, the lessons on Polarization found in Section 4, and it ought to be accompanied by the reading of Section 6, on Interferometry and Relativity. The Sagnac interferometer is also the natural tool for doing the Electro-Optic experiments of Section 15.

- Section 7, on the Mach-Zehnder interferometer, is another alternative subset. With enough polarization skills from Section 4, both kinds of Mach-Zehnder interferometers can be built. Given a polarized-light Mach-Zehnder interferometer, it's very natural to read section 8, on Interferometry and Quantum Mechanics.

The Experiments in Sections 9-16 are all more or less independent of each other, and any subset of them might be done; but each of them does depend on having one or another kind of interferometer assembled.

- As mentioned above, Section 15, on the electro-optic effect, assumes a working Sagnac interferometer.
- Sections 9, 11, and 12 can be performed using just a basic Michelson interferometer, though each is easier using a quadrature Michelson interferometer.
- Sections 10, 13, 14, and 16 really require the quadrature Michelson, and this interferometer also makes the previous interferometric measurements easier and more illuminating.
- Section 11 on index of refraction can be performed using the Michelson, Sagnac, or the Mach-Zehnder interferometer.

Appendix B. Laser safety

Laser beams deliver energy to the human eye and can damage it permanently, so it is important to understand the safety precautions required. For the case of low-power, continuous-wave, visible laser beams as in this Interferometry apparatus, the safety requirements are easy to meet with a trifle of common sense. In a simple pair of lessons:

- it is NOT SAFE to let a collimated laser beam fall into your unprotected eye, even very briefly;
- but laser light scattered from a diffuse surface IS SAFE to look at, even from close range and even for long times.

Here's the support for those two claims, taken from the Z136.1 laser safety standard of the Laser Institute of America.

We work out the worst-case scenario, in which a 5-mW diode laser beam delivers its light into an area about 1 mm x 2 mm, or 2 mm^2 = 0.02 cm^2. This gives a power density of 5 mW/0.02 cm^2 = 250 mW/cm^2, which is more than double the power density of full sunlight. Just as you're smart enough not to look at the sun, or let its full irradiance get focused onto your retina, so you need to avoid looking into a laser beam.

The contrasting case of scattered light is very different. For light falling on a diffusing surface, like the bead-blasted and anodized aluminum surfaces of your optical mounts, or the powder-coated surfaces of your alignment towers and alignment paddles, it is reasonable to assume the light scatters into about 2π steradians of solid angle. Suppose a 5 mW laser-beam's spot on such a surface is viewed from a mere 10 cm = 4" away; even if 100% of the light is scattered, the irradiance at the eye is only 5 mW/628 cm^2 = 0.008 mW/cm^2 = 8 $\mu W/cm^2$.

Here for contrast are some safety limits taken from the LIA Z136.1 standard:

- For exposure times shorter than 18 μs, the safety limit is 5×10^{-7} J/cm^2; this is to be compared with worst-case delivered energy of (250 mW/cm^2) (18 μs) = 4500 nJ/cm^2 = 45×10^{-7} J/cm^2, which is *nine times higher* than the safety limit. You may thereby conclude that even very brief experiences of the full laser beam falling into your eye are <u>unsafe</u> by this standard; that's why lasers ought to be aimed safely first, and turned on only later, and why every laser beam ought always to end at a safe destination (such as an absorbing or a diffusing screen, NOT a mirror or other specularly-reflecting surface).
- By contrast, for exposure times as long as a full 8-hour day, the safety limit is given by C_B $\mu W/cm^2$, where $C_B = 10^{15(\lambda - .550\ \mu m)}$. This gives a safety limit of 18 $\mu W/cm^2$ for 633-nm laser light, so the scattered irradiance of 8 $\mu W/cm^2$ computed above lies *below half* of the safety limit.

Appendix C. Laser parameters

You may want to know some technical data on your laser sources:

a) HeNe laser
Manufacturer is JDS UniPhase, model number is 1107P
Output has wavelength (in standard air) 632.8 nm, power > 0.8 mW
Power stability is ±5%, due to 'mode sweeping'
Output is >95% TEM_{00} mode
Polarization is >500:1 linear polarization;
E-field vertical when laser head's power cord is rotated to low point of case's periphery
TEM_{00} mode has a beam diameter (full width between $1/e^2$ points) of 0.48 mm,
and a beam divergence (full angle between $1/e^2$ directions) of 1.8 mrad
Nominal longitudinal mode spacing is 1090 MHz

Note that in addition to the white-paper peel-off tab that covers the output port of the HeNe laser for shipping, there are two other output safety features:

1) There is a few seconds' time delay between turning on the power and the appearance of a beam, giving you time to change you mind if you realize your eye is in an unsafe place relative to the laser.
2) On the output-end faceplate of the laser, there is a screwdriver slot; if it's rotated through 90 degrees, a mechanical shutter blocks the output beam.

b) Diode laser
Manufacturer is US-Lasers, Inc., model number is M650-5
Output has wavelength 650 (-5, +10) nm nominal, power 5 mW nominal
Power stability is 0.1%, due to APC = automatic power control
Output is single transverse mode, but *not* cylindrical
Polarization is >100:1 linear polarization;
orientation of E-field can be adjusted by rotating brass laser housing within aluminum support structure (loosen one set screw)
Divergence of output depends on focusing,
adjusted via rotation of inset lens (via shallow screwdriver slot in front of brass housing)

Appendix D. Vibration control via the stiffening ribs

Interferometry is a technique famously sensitive to small mechanical displacements of optical components, and presents a technical challenge to the mechanical stability of the apparatus used to perform it. You'll recall that interferometric signals are sensitive to sub-wavelength displacements of mirrors, so that 0.3 μm is a *large* displacement by interferometric standards; meanwhile, the apparatus can easily have a scale of 300 mm or 300,000 μm, so that part-per-million changes in position matter significantly. So how is the required stability achieved?

You've seen previous mention of the need for temperature stability and uniformity, and the need to prevent inhomogeneous packets of air from blowing though your apparatus; this section considers the more immediate threat of vibrational displacements of the parts of an interferometer. There are two quite distinct lines of approach here:

First, one can imagine isolation of the interferometer from external vibration. Apart from the vibrations of sound (ie. coupled through the air) there are vibrations of the base (ie. coupled through whatever supports the interferometer against falling). There are vibration-isolating tables of various levels of sophistication to which you might have access, and they can make moderate to spectacular improvements against typical building vibrations.

Second, one can imagine that an interferometer might still vibrate, but can think of ways of stiffening it so that all its parts vibrate innocuously as a single rigid unit. This line of thought has motivated the design of the very stiff TeachSpin mounts for mirrors and beamsplitters, and it further suggests that bending of the relatively thin aluminum optical breadboard is the largest remaining source of non-rigidity in the apparatus.

If you've built a Michelson interferometer on the breadboard, and supported that breadboard on an ordinary (ie. not a vibration-isolating) table, you've seen the consequences. As an illustration of the importance of rigidity, it would now be worth installing the simple stiffening ribs that are designed to be bolted to the underside of your aluminum breadboard.

The figure below shows the ribs in their intended positions; you'll note that the 1/4-20 socket-head cap screws sent along with the ribs are deliberately of short length, so as *not* to use up the whole depth of the threaded holes in the breadboard. This keeps the holes useable for optical mounting from the top of the breadboard.

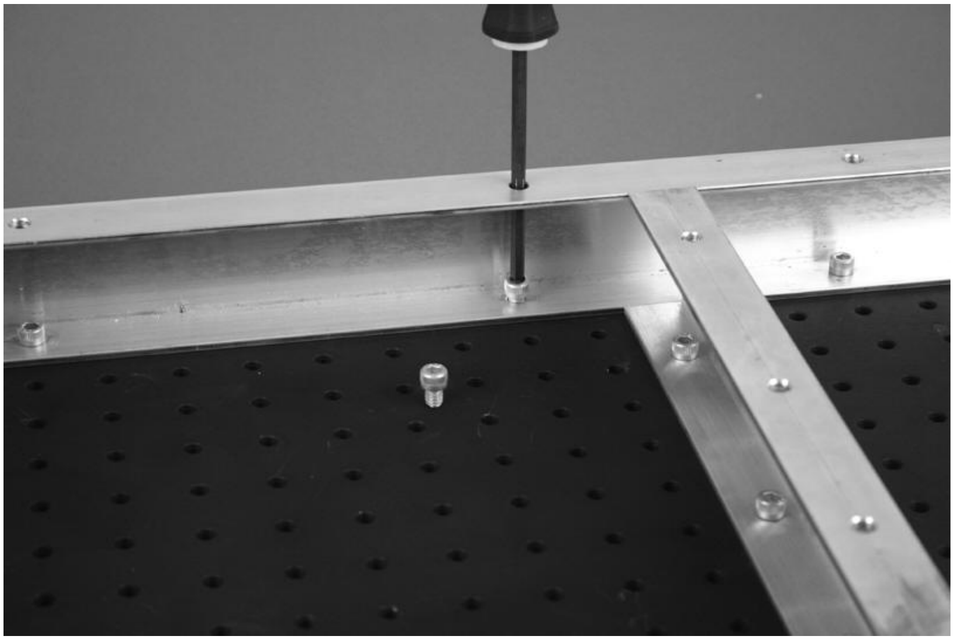

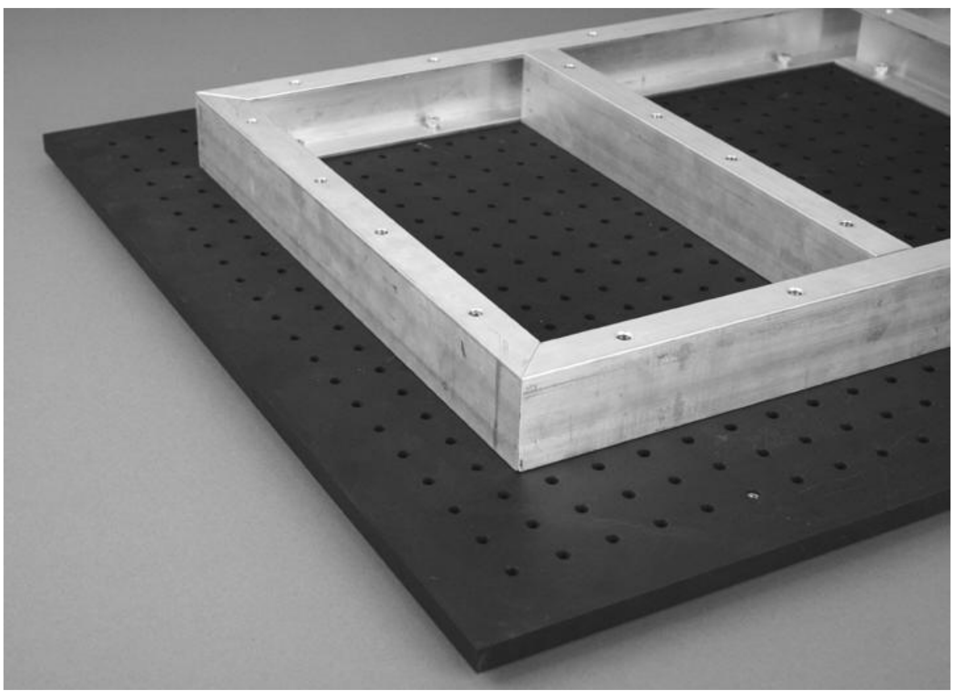

Figure D-1: The intended method for mounting the stiffening ribs

Once you have the ribs installed, you are also welcome to try supporting the ribbed structure from below by its four corners on the four 2 x 2" Soft urethane pads, which are designed to damp higher-frequency vibrations from coupling up to the breadboard. Of course you can try your own recipe for further vibration isolation, but you should already be able to see a marked improvement in the stability of your interferometric signals.

Appendix E. The 'bull's-eye' fringes from a Michelson interferometer

Especially when viewing fringes from a Michelson interferometer illuminated by an angularly expanded beam, it is easy to see a series of circular bright fringes forming a memorable 'bull's-eye' pattern. Since this is rather different than the parallel straight fringes seen in the Sagnac interferometer, it is worth working out the differences between the two patterns, especially because some useful information can be extracted from the bull's-eye pattern.

The typical way to form an expanded region of illumination on a viewing screen is to turn the input collimated laser beam into a divergent beam using a lens, and this is the key to the difference between the straight-line and bulls-eye fringe patterns. So now we need to treat the wavefronts of the waves as *spherical* surfaces. The second key ingredient in explaining the pattern is to realize that in an unequal-arm Michelson interferometer, the two overlapping beams represent wavefronts that have come different distances from a common source. Thus the spherical wavefronts that interfere are now going to have values for their radii of curvature that *differ* (by twice the one-way arm-length difference in your interferometer). So the bull's-eye pattern is not so much an diagnostic of angular misalignment as it is of arm-length difference. (Hence in the Sagnac interferometer, where the two interfering paths have identical length, there's no reason to expect a bull's-eye effect. Similarly, even in an unequal-arm Michelson interferometer illuminated by a transversely expanded beam with *plane* wavefronts, the bull's-eye pattern is not expected to occur.)

Now suppose that two spherical wavefronts, emerging from virtual point sources on the same x-axis, and of respective radius of curvature R_1 and R_2 where they overlap, are mutually tangent and constructively interfering at a point where the x-axis pierces a viewing screen. The picture diagram below (not to scale) suggests the mechanism for the bulls-eye pattern: as one goes off-axis on the screen, the wavefronts are no longer overlapping in the way needed to keep the interference constructive. If the center of the interference pattern is a bright spot, then the first dark fringe will form a circle around that spot when the phase difference is π, or when the appropriate distance difference is $\lambda/2$. As we've supposed the center of the pattern to be a bright spot, the nth bright ring will occur when the appropriate distance difference is $n\lambda$.

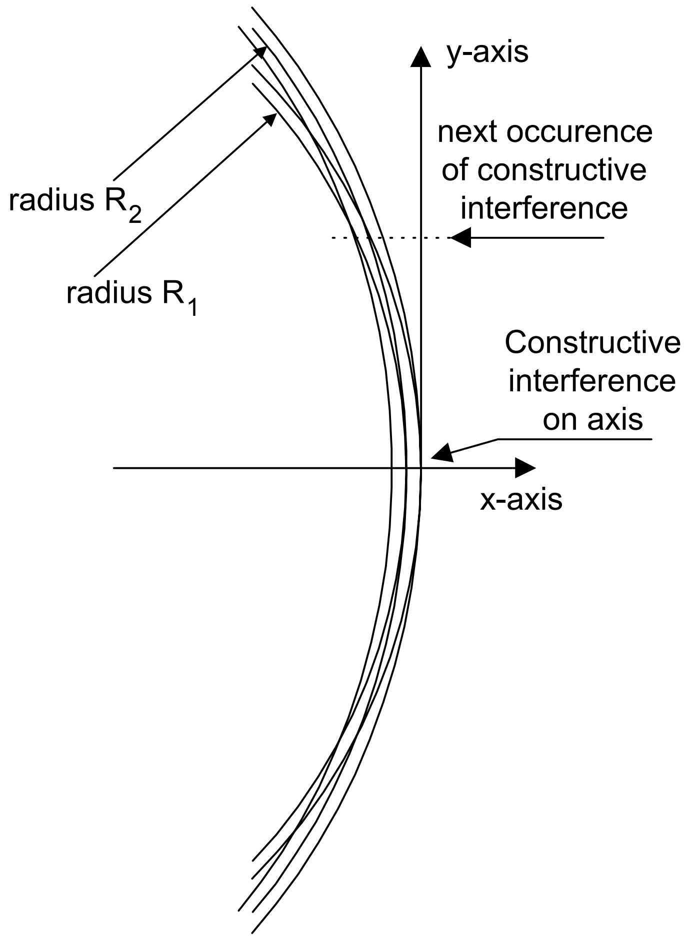

Figure E-1: Interference between two sets of spherical wavefronts (of distinct radii of curvature) showing the mechanism for producing the 'bull's-eye' pattern

The geometry best suited to the distance calculation is shown below. A circular wavefront tangent to the x=0 plane, emerging from a source at x = -R, is given by the locus

$$(x+R)^2 + y^2 = R^2 \text{ ; or } y^2 = R^2 - (x+R)^2 = 2R(-x) - x^2 \approx 2R(-x) \quad .$$

For a given distance y off-axis, the wavefront has pulled away from the x=0 tangent plane to a bent-back position

$$x = -\frac{y^2}{2R} \quad ,$$

and two wavefronts, of differing radius of curvature R_1 and R_2, will interfere constructively at off-axis distance y_n, occurring when

$$\Delta x = \left(-\frac{y_n^2}{2R_2}\right) - \left(-\frac{y_n^2}{2R_1}\right) = n\lambda \quad .$$

This gives the radial distance to the nth bright ring as

$$y_n^2 = 2\, n\, \lambda \frac{R_1\, R_2}{R_2 - R_1} \quad .$$

This predicts that the rings of the bull's-eye will *not* be equally spaced, but instead have

$$y_n = y_1 \sqrt{n}\ , \text{ where } y_1 = \sqrt{2\lambda \frac{R_1 R_2}{R_2 - R_1}}\ .$$

In practice, the distances R_1 and R_2 might be about 0.5 m, while the difference R_1 - R_2 can be taken to be double the one-way path difference in the Michelson interferometer, perhaps 2(0.025 m) for a one-inch displacement of one mirror from the equal-arm condition. This predicts are first bright ring of radius about 2.5 mm around a bright central spot, and a characteristically diminishing spacing between rings farther out in the ring pattern.

This results also predicts the bull's-eye pattern fails to exist when $R_1 = R_2$, as appropriate to the equal-arm Michelson interferometer. From such an interferometer, illuminated by a diverging beam, there will instead emerge on a viewing screen a fringe pattern of unpredictable shape, diagnostic of the imperfect flatness of the two end mirrors used in building the interferometer.

Appendix F.   The flexure stage and its mechanical train

The flexure stage included in the TeachSpin kit is intended to be an alternative to the 1"-thick baseplates which support the mirror uprights. Its distinctive feature is the freedom it gives to a mirror to translate by ±1 mm relative to its central position. The challenge is to permit this translation, while suppressing other translations or rotations that the mirror might undergo. This appendix discusses first how the flexure is built, secondly how it's intended to be used, and finally how the differential micrometer performs its magic.

a.   Flexures in general and the 8-bar flexure in particular

There are lots of ways to achieve the relative motion of structures, and you might be familiar with ball bearings (for rotational motion) and sleeve bearings (for translational motion). But in the context of interferometry, there is considerable incentive to avoid all forms of rolling or sliding, because of issues of grip-and-slip or 'stiction'. One alternative to these is the systematic exploitation of the bending of elastic metals, via the use of metallic links that 'flex' during the intended operation. You'll note that elastic deformation involves (in principle) only the stretching, not the breaking, of inter-atomic bonds, and that (in principle) it offers a reversible and stiction-free way to allow motion.

The simplest flexural element is a 'flexure hinge' as illustrated below, which is easily built out of a solid single piece of metal, and which can be very stiff indeed against deformation in any degree of freedom other than the desired bending sort of motion.

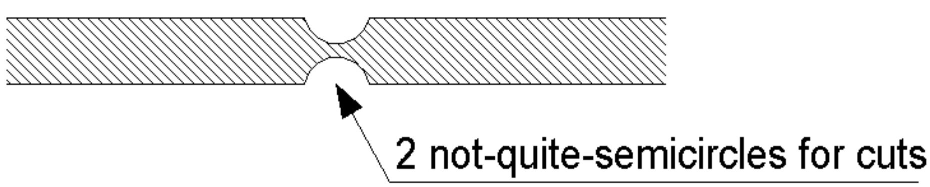

Figure F-1: A flexure hinge formed by machining solid metal

But a single hinge (as used in the TeachSpin mirror mounts) gives a motion more like rotation than translation, so the next step toward understanding your flexure stage is to consider the 4-bar, 8-hinge structure that can be formed out of simple flexure hinges. Next to the flexural element is shown a more humble analog, imagined to be built using two wooden bowstaves.

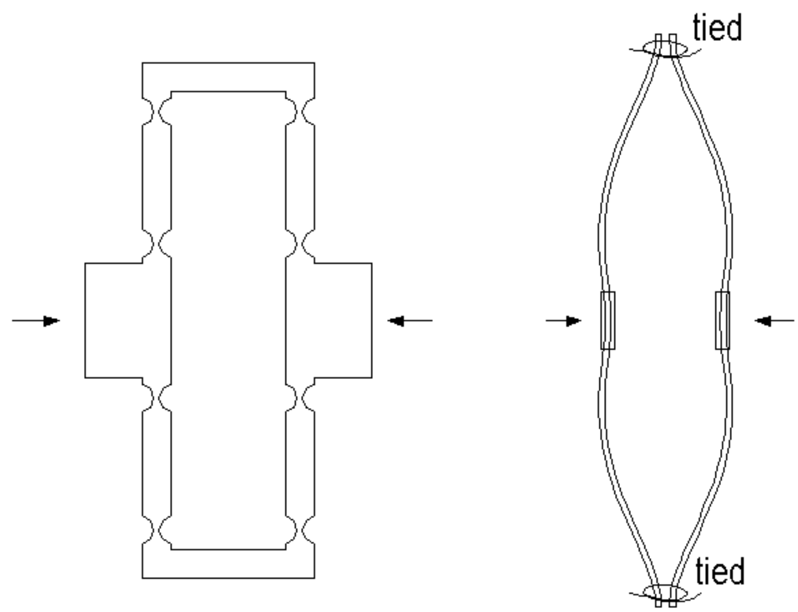

Figure F-2: A four-bar hinge, and equivalent made of tied bowstaves

Now if you imagine two persons each grabbing a bowstaff by its grip, you can easily imagine the elastic deformations that will allow them to change the spacing of the two handgrips -- that's shown by the arrows in the diagram.

You can understand the analogous motion in the 4-bar flexural element most easily by imagining that one of its attachment points is bolted to a rigid wall, while your hand is gripping the other attachment point. Clearly you can pull away from, or push toward, the wall, giving the desired degree of translational freedom. Given symmetry of construction above and below the centerline, there will be no rotation accompanying this translation. And perhaps you can imagine the considerable resistance the structure would present to the two other (perpendicular) attempted translations of the attachment point.

Your TeachSpin translation stage uses *two* of the sub-systems you've just come to understand, as illustrated in the figure below -- first with bowstaves, then with hinge flexures:

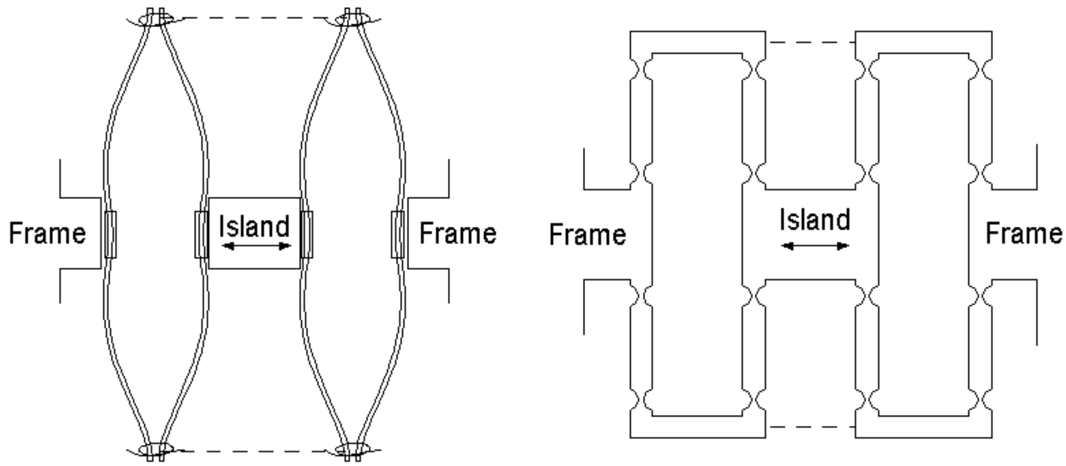

Figure F-3: Supporting a central island with double 4-bar hinges; first using bowstaves, and then the equivalent built from metal

Now you can see a central 'island' or 'stage' that can move with the desired translational degree of freedom, while supported at both its left and right ends by 4-bar flexural elements. You can see that the two outermost attachment points can both be connected to a single, external, rigid frame. Finally, you can imagine that the tied tips of the bowstaves can have various undesired motions suppressed by the addition of two more rigid rods, shown in the bowstaff diagram as dotted lines.

Now if you handle your one-piece TeachSpin flexure translation stage when it's free of the optical breadboard and free of a mirror-mount upright, you'll be able to see all these elements present and working together. If you hold the external frame with one hand, and grip the central stage with the other, you'll be able to feel the relative motion (of the stage relative to the frame) that it permits. You are also free to try to get the stage to move in any *other* direction of translation or rotation relative to the frame, and to feel and intuit how very stiff it can be in these other degrees of freedom.

With any flexural element, there is the need to keep the deformation of the bending structures within the elastic limit of the material in question. In the TeachSpin translation stage, the motion is kept within these limits via four 'limit screws' mounted in the frame. These 6-32 socket-head setscrews have been adjusted to limit the motion to ±1 mm relative to the central position, and these limits should not be exceeded.

b. Controlling the motions of the flexure stage

If you mount a mirror-bearing upright on the flexure base, and build that mirror into a simple Michelson interferometer, you will have trouble with the visual detection of any fringes at all. If you try electronic detection instead, you might be able to accomplish an alignment that results in fringe signals, but they might look something like this:

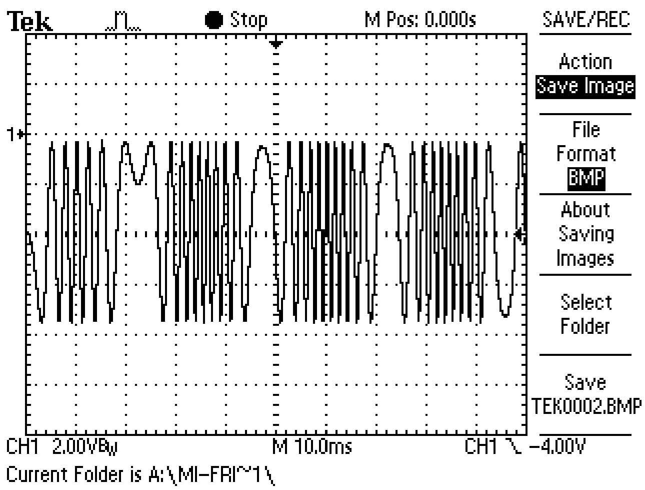

Figure F-4: Interferometric fringe signals due to an end mirror, freely vibrating on the unrestrained flexure-translation stage

There is a *lot* of information in such a signal, including both the amplitude and the frequency of a mirror motion (and how do you *extract* this information?), but for the moment what it's telling you is that the unconstrained 'island' of the translation stage is vibrating rather freely and with excessive amplitude. So in real operation, the stage needs to be constrained to lie at a particular point, and not vibrate through a range, of its possible translational motion.

The TeachSpin design is to use a (hidden) push-spring to drive the island toward one extreme of its possible range of motion, and then to use an externally driven pushrod to push the stage back against the spring. This is shown schematically in the drawing below:

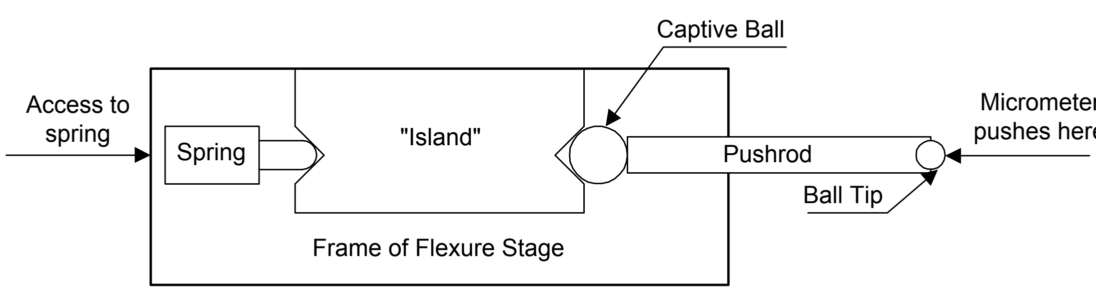

Figure F-5: The locations of the push-spring and the pushrod

Here the push-spring is on the left, and pushes the stage to the right (with a force of order 30 Newtons). The amount of force can be adjusted, by rotating the whole spring-bearing push-spring assembly within its threaded holder inside the frame of the translation stage (use a 5/64" ball-tipped Allen driver to gain access to this adjustment, and for fun, start by withdrawing the whole push-spring assembly altogether to inspect it).

The push-spring might be (re-installed and) adjusted to push the 'island' a half or full millimeter to the right of center; then the pushrod, located on the right and pushing back toward the left, will have control over the position of the island and the mirror mounted on it. The simplest way to gain this control is to have the differential micrometer push, to the left, on the right-hand end of the pushrod; then translations of the tip of the micrometer will be transmitted directly to the motion of the mirror.

c. How a micrometer, and a differential micrometer, work

The micron, or mi'-cro-me-ter, is a unit of length, but a mi-cro'-me-ter is a machinist's tool. This section explains how an ordinary micrometer 'head' allows precision control of translation, and how a differential micrometer offers higher resolution still.

There's a micrometer head present and working in the 0-1" translation stage that's part of the TeachSpin kit, and to understand its operation, all you need is a nut and a bolt.

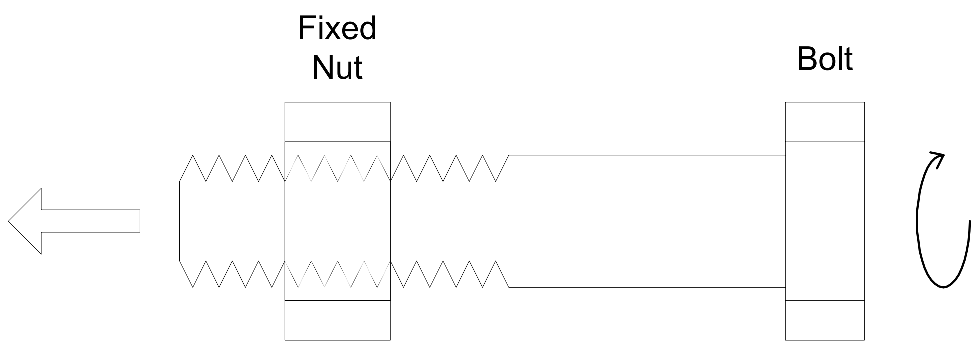

Figure F-6: A nut and bolt to illustrate an ordinary micrometer

Now imagine that you hold the nut fixed in space, and rotate the head of the bolt. After one full turn, the whole bolt will have translated in space by the pitch of the screw. [For example, for a 1/4-20 nut and bolt, the pitch is 1/20" = 0.050" = 50 mils = 1270. μm.] What's more, if you can execute one-tenth of a full turn of the bolt, you'll get one-tenth of this translational displacement.

In your 'ordinary' 0-1" micrometer, the barrel that you handle is an expanded version of the head of the bolt, and the screw thread in question is a 40-pitch screw. That's why one full turn of the barrel advances the micrometer's tip by 1/40" = 0.025" = 25 mils = 635 μm. That's also is why the barrel is labeled with 25 marks around its circumference, to permit the easy delivery of incremental rotations of one twenty-fifth of a turn, and hence translations of 0.001" = 1 mil = 25.4 μm.

Finally, along the non-rotating body of the micrometer there are marks every 0.025", to allow you to keep track of integral numbers of turns; hence you read a micrometer coarsely along this scale with markings every 25 mils, and finer subdivisions with the 1 mil markings on the barrel itself.

Two last issues with an 'ordinary' micrometer: It's far from trivial to make a screw with a pitch much finer than this (80- and 100-pitch screws are right at the state of the art) so there's not too much more resolution to be gained; and secondly, you'll note that the working shaft of the micrometer rotates while it translates, which is undesirable in some applications.

So now to the 'differential micrometer' that solves both problems, allowing higher positional resolution and delivering pure translation of its working end. To understand how it works, you'll want to think back to a bolt, but this time where the bolt is free to translate in a sleeve but not to rotate at all. For concreteness, think of a 1/4-20 bolt with a pitch of 1/20" = 0.0500". Now imagine that a nut, constrained between two plates, can be made to rotate, and you'll deduce a translation-only motion of the bolt, by 0.05" per turn of the nut:

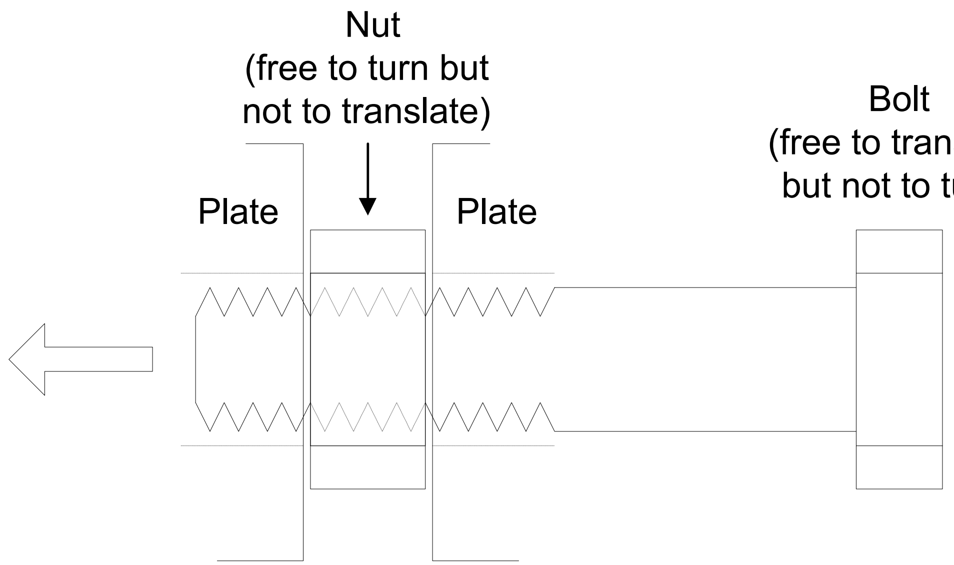

Figure F-7: A nut constrained between two plates, with a bolt constrained **not** to be able to rotate

This solves the rotating-shaft problem, but not the resolution problem. So now imagine that the walls confining the nut are removed, and that the length of the nut is extended, and (crucially) that the *outer* surface of the nut is no longer a hexagonal shape, but is itself round and threaded -- perhaps with a 19-pitch screw. Now imagine an outer fixed sleeve, with an internal 19-pitch screw, for the long nut to be threaded into.

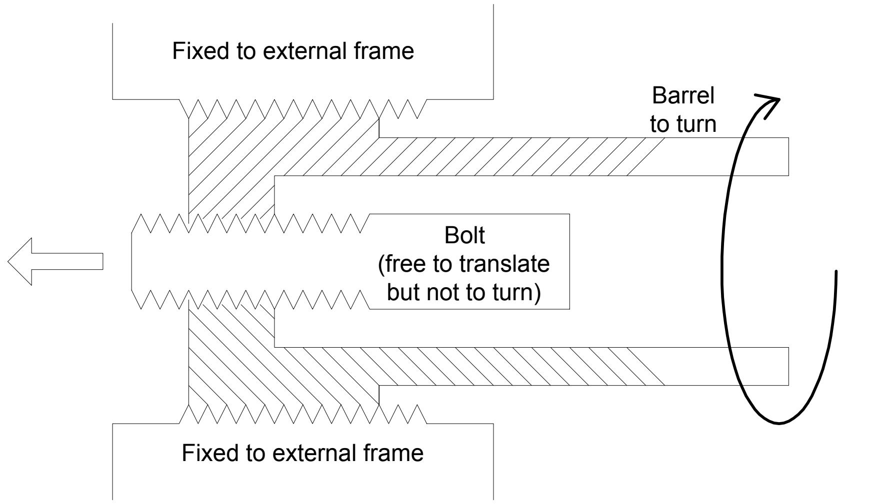

Figure F-8: A nut with internal **and** external threads, mounted in a fixed outer threaded housing, with a bolt constrained **not** to be able to rotate

So now think what happens when the outer frame is fixed (no rotation, no translation) and the barrel is rotated by one full turn. The whole barrel moves (to the left, say) by 1/19" = 0.05263..." = 52.63... mils. But the nut is *also* rotating one full turn relative to the central 'bolt', since the bolt is constrained so as not to be able to turn. It follows that the bolt will have to translate, by 1/20" = 0.05000" = 50.00 mils, relative to the nut.

If both the 19- and the 20-pitch screws are of the same handedness, then the net displacement of the bolt relative to the external sleeve is the *difference* of these two translations; in this example, it's (52.63... - 50.00) = 2.63... mils. Hence the name 'differential micrometer': without having to try to build a screw with this tiny pitch, we nevertheless get this very small and controllable translation of the working tip of the micrometer for each turn of the 'nut'.

Finally to the Mitutoyo differential micrometer you're actually using: it does indeed have a non-rotating shaft (the 'bolt' of the figures above) and it does have a rotating 'nut' (which is the barrel you actually touch and turn). That 'nut' has hidden internal and external threads of slightly different pitch, such that the shaft advances by just 50 μm (= 1.968... mils) per turn of the barrel.

As with an ordinary micrometer, so in this case too, the barrel's periphery is subdivided by marks; here each one corresponds to one fiftieth of a turn of the barrel, and hence each one stands for a translation of 1 μm. For a final touch, there's a Vernier scale on the frame of the micrometer, against which you can compare these subdivisions, to allow interpolation to one-tenth of a smallest barrel division, or 0.1 μm of translation.

It is quite remarkable that a mechanical object can be made which allows a full range of 2.5 mm = 2500 μm of translation, but yet allows readability to a resolution of 0.1 μm = 100 nm. More remarkable still is the silky feel of all those contacting threads, and the entire freedom from backlash in the motion of the micrometer's tip.

Appendix G. 'Shimming' the optical mounts

The TeachSpin beamsplitter and mirror mounts are built for maximal rigidity, and to achieve this they have a minimal number of degrees of freedom. The beamsplitters contain no adjustments at all, and the mirror mounts have only one rotational degree of freedom. But that entails that there will be times in interferometer alignment when you seem to lack a knob to achieve the desired adjustment. Here's how you can gain the needed degrees of freedom of adjustment, with*out* any sacrifice of the rigidity of the mounts. The technique is called 'shimming'.

We start with a concrete example, of a single input beam headed horizontally across the optical breadboard, and a beamsplitter (either of the flat-plate, or polarizing-cube, variety). A fraction of the input beam continues right through the beamsplitter, and its angular position is still under control of steering mirrors upstream. The question is how to adjust the direction of the *deflected* beam coming from the beamsplitter. Nominally, that new beam is also propagating horizontally, and has been deflected by exactly 90°, but in practice you might need some way to trim up both those angles. Here's how to do the shimming -- first in one direction:

> The nominal 90° deflection depends on the rotation, about a vertical axis, of the whole beamsplitter mount on the breadboard. To exercise this rotation, you need to loosen the two 1/4-20 screws that clamp the mount to the table, and use the freedom provided by the somewhat-oversize clearance holes in the base of the mount. Use two alignment towers to monitor the state of the transmitted, and of the deflected, beams as you rotate the mount between the limits still provided by the screws. You ought to see that the transmitted beam is unaffected, but the deflected beam is steered, by this adjustment. When you have the deflected beam going in the direction you wish, then you can carefully snug down the two mounting screws, fixing the mount in the orientation you need. The adjustment you've achieved is not as smooth and continuous as that provided by a thumbscrew, but in return, the shimmed mount is as rigidly attached as it can be.

Now back to your view of the deflected beam on its beam-tower: you've been able to translate its impact point horizontally, side-to-side, but what about moving that spot around vertically? Here's another sort of shimming adjustment -- for vertical angle.

> Go back to loosening the two mounting screws of the mount to the breadboard, and now instead of rotating the mount around a vertical axis, try lifting one of the corners of the base of the mount vertically off the breadboard. If you've loosened the screws by two full turns, you could lift any corner by about 0.1" or 2.5 mm. Have a look, on your alignment towers, to see what that does for the state of the transmitted, and the deflected, beams. You should again see negligible effects on the transmitted beam. But now, if you've found the correct corner to lift, you should see a new <u>vertical</u> degree-of-freedom of the state of the deflected beam.

Once you've identified this effect, you know which corner needs to be lifted. To get a smoother way to achieve this, look onto the top of the base of the beamsplitter mount, and find the two 10-32 threaded holes located right near the mounting holes you've been using. One of these two holes ought already to contain a swivel-headed 10-32 setscrew; use a 3/32" Allen wrench to back it all the way out for your inspection. Now install that setscrew into the threaded hole near the corner you need to lift, and drive it downwards with your Allen key until its swivel head contacts the breadboard. You should be able to see that further driving starts to lift that corner of the mount up off the breadboard, and as a result steers your deflected beam in the direction you desire. You now have a pretty smooth way to adjust the mirror mount to the degree of shimming you require.

Of course, the two 1/4-20 mounting screws have to be loosened for this to work at all, and of course, that spoils the other degree of freedom you previously exercised. Furthermore, you should appreciate that a beamsplitter mount, one corner of which is propped up on a swivel-headed setscrew, is not so rigidly attached to the breadboard, as you might desire. So here's a way to regain the rigidity you want:

> After you've used the swivel-headed screw to determine how high you need to lift a corner (remember, each full turn after tabletop contact buys you 1/32" = 0.031") you can prepare a 'shim' of the appropriate thickness that will replace the effect of the setscrew. If you need 0.008" of lift (ie. about a quarter-turn of the setscrew), you can cut (with ordinary scissors) some strips of the stainless-steel shim stock provided with your kit. The pieces you'll find have nominal thicknesses of 0.002" and 0.006", and you can stack the pieces you require. You only need shim sizes of about 1 cm by 3 cm to fit under the appropriate corner of the beamsplitter base to lift it by the amount required.

The figure below shows a beamsplitter base with some shim stock under one of its four corners.

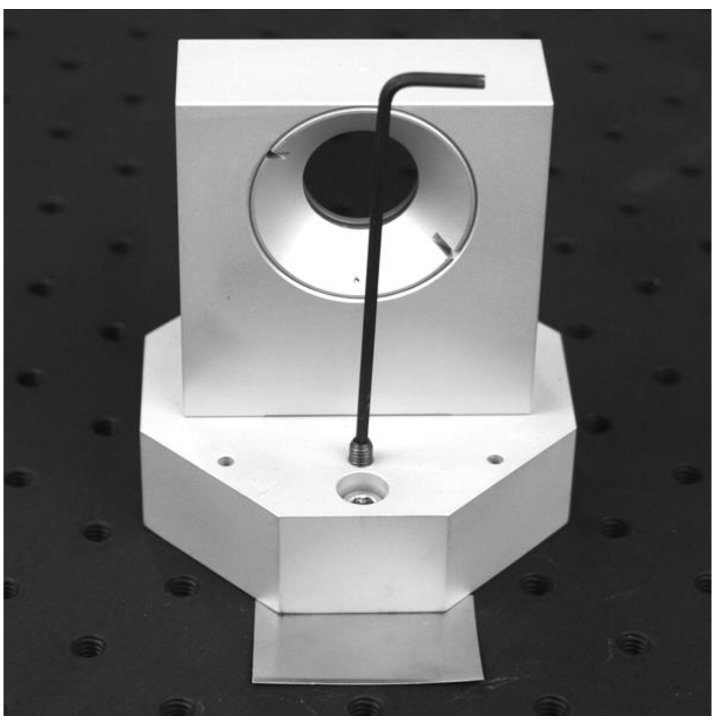

Figure G-1: Using shim stock to elevate one corner of a beamsplitter's mount; also shown is position of elevation set screw with swivel base

Once the shim stock is in place, you can back off the swivel-headed setscrew, and you can still exercise the previous rotational degree of freedom of the whole mount, before finally tightening down both mounting screws for a mounting nearly as rigid as could be achieved directly to the table.

This procedure sounds a bit complicated, but it can be rapidly learned and applied to each new interferometric set-up that involves the placement of a beamsplitter. The same techniques of 'shimming', in rotation and tilt, can be applied to the mirror mounts too -- here's how.

Suppose you have a mirror mount with a horizontal hinge line, so that the beam reflected from it already has a built-in thumbscrew adjustment of a laser beam's vertical position on an alignment tower or adjustment paddle. How do you get the necessary *horizontal* adjustment? By loosening the mount's table mounting screws a bit, and rotating the whole mount about a vertical axis on the breadboard, and then re-tightening the screws.

Suppose you have a mirror mount with a vertical hinge line, so that the beam reflected from it already has a built-in thumbscrew adjustment of a laser beam's horizontal position on an alignment tower or adjustment paddle. How do you get the necessary *vertical* adjustment? By loosening the mounts table mounting screws a bit, and lifting either the front or the back edge of the base above the tabletop by a trifle. (Here there's no swivel-headed setscrew to employ.) Once you've identified the edge you need to lift, you can slip under it the right thickness of shim stock, and clamp back down for a truly rigid but now properly adjusted attachment to the optical breadboard.

You'll note that a 0.002" piece of shim stock, emplaced to create a tilt on a baseline of about 3", will give a tilt of about 2/3 milliradian to a beamsplitter or mirror mount. This in turn will steer a reflected beam by about 4/3 mrad, which is somewhat smaller that the divergence already present in the beam. Typically you will not need finer resolution than this in the shimming operation.

Not every interferometer requires all this shimming; many times you can get everything adjusted correctly with no shim stock at all. But it is good to know that you can retain the full rigidity of the TeachSpin mounts, meanwhile exercising (where necessary) some degrees of freedom that are not provided for by thumbscrew adjustments.

Appendix H. The mounts for 1" x 2 mm beamsplitter plates

Your first interferometer will probably be a simple Michelson interferometer, in which a 50/50 beamsplitter needs to be held in the incident beam at a 45° angle of incidence. This is accomplished using a beamsplitter holder, mounted on an octagonal base, attached to which is an 'upright' designed for holding thin flat optical substrates of 25 mm or 1" diameter, having thickness of about 2 mm.

[When using polarizing beamsplitter cubes (PBSCs) for the same function, you'll use the same bases, but different uprights -- read about their use in Appendix N.]

There are three optical elements in your kit that might be mounted in these mounts: they are the 50/50 dielectric-coated beamsplitters, a metal-film coated beamsplitter, and a compensator plate. All of these are built on BK-7 glass substrates, and all have their inactive surfaces anti-reflection coated. [See Appendix R for the meaning of BK-7.] The dielectric-film beamsplitters are identifiable by the yellow or blue sheen they seem to have in reflection; by contrast, the metal-film beamsplitter has a pronounced grayish cast, and the compensator plate is nearly invisible.

The design of the beamsplitter mount is shown (in top view) below. The optical element is held into place by a conical threaded insert, which bears on the optical element via a nylon washer of about 0.90" internal, and 1.98" external diameter. Find the black fiberglass tool of width 1.96" that serves as the 'spanner' or wrench by which the conical insert is to be threaded into or out of the beamsplitter upright, and use it to loosen or tighten the threaded insert into place.

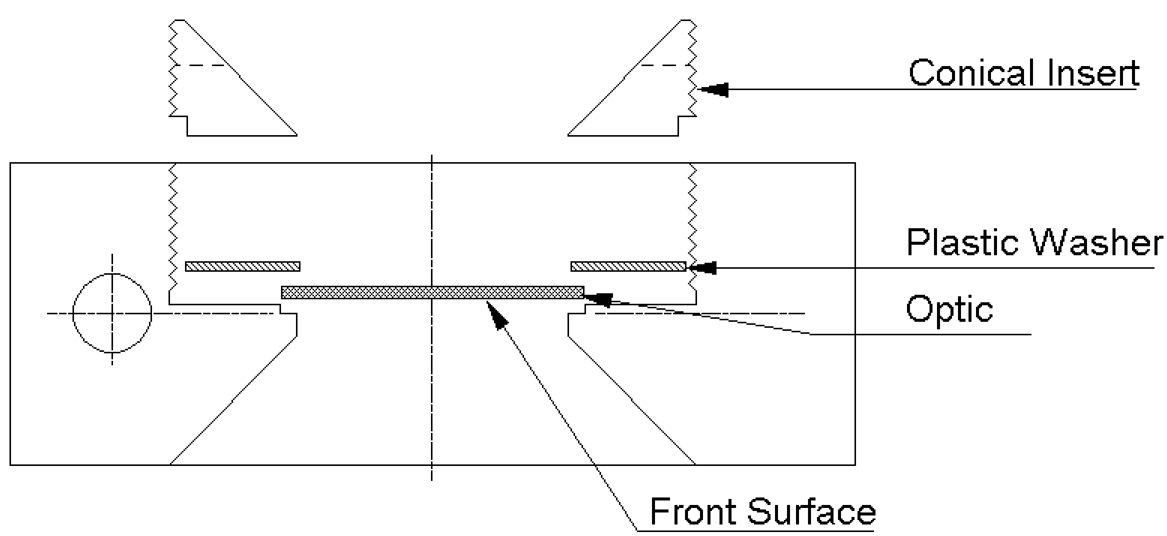

Figure H-1: A horizontally sectioned drawing of the beamsplitter-mount's upright, showing the beamsplitter plate, the washer and the conical insert

The installation or removal of a beamsplitter is best conducted with the whole mount removed from the optical breadboard, and laid out with the optical surface approximately horizontal. Now turn it, as in the figure below, so the threaded insert is on top, and remove the insert with your fiberglass tool. You'll get a view of the nylon washer, which you can remove (carefully!) with a plastic tweezers. That will expose the optic itself, lying centered in a shallow depression of 1.00" diameter and 0.025" depth. If you need to remove the optic, do so with plastic tweezers. The goal, or course, is to handle the actual optical element by its edges only.

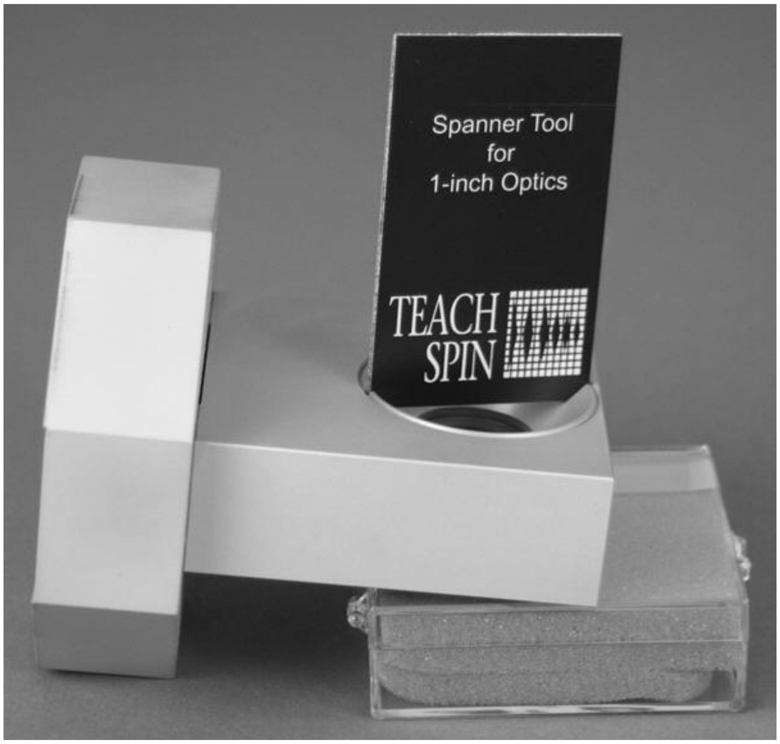

Figure H-2: The spanner tool for access to beamsplitter plate (shown being used in the recommended orientation)

The design of the beamsplitter mount is to have the coated working surface of the beamsplitter lie face *down* in the mount as you're holding it. That will place the actual beamsplitting surface directly above the reference hole in the optical breadboard, and it will put the thickness of glass substrate on the same side of the beamsplitting surface as is located the conical insert. [This all matters, a lot, when you're trying to do white-light interferometry.] So now you need a way to tell which face is which of the optical elements.

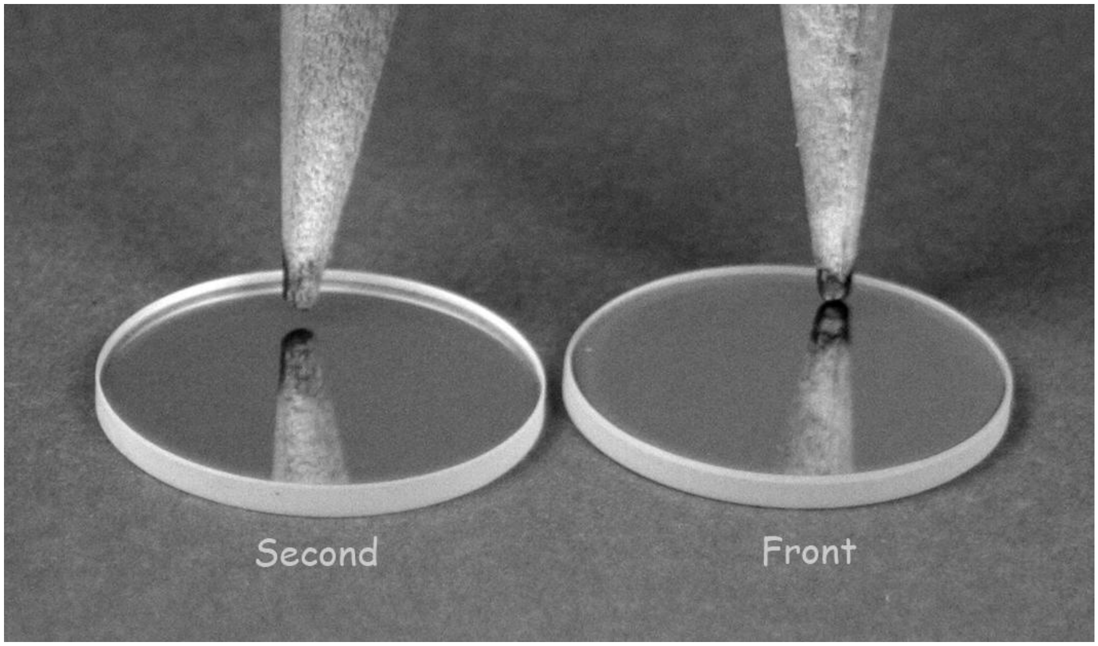

Figure H-3: Oblique views of two 50/50 beamsplitter plates with contacting objects to show the difference between ‘second-surface' and ‘front-surface’ reflections

- The 50/50 beamsplitters have a dielectric film on their working faces, and above is a view of two of them at oblique incidence, each being touched with the tips of ordinary pencils. [The goal is to touch it very gently, with a non-scratching tool, just for this eyeball test.] You can see the difference between the 'second-surface reflection' and the 'front-surface reflection' in the two views, and thereby identify for certain the face with the dielectric coating.
- A similar test is feasible, but much harder, for the metal-film beamsplitter. (You'd have to look more nearly at grazing incidence.) So instead, you may inspect the 2-mm-wide ground-glass edge of the metal-film beamsplitter, and scan around its periphery for a little pencil mark of the shape < or > written onto the glass. That's an arrow, and it points to the face having the metal-film coating.
- The compensator plate has *both* its surfaces anti-reflection-coated, so there's no distinguishing them. You'll probably only be mounting this in the 1" lens-holder, when using the plate in the rotation stage, for white-light interferometry.

Remember when you mount a beamsplitter into the upright, the goal is to have the working face seated down into the upright's shallow depression, and to have the nylon washer and threaded conical insert bear on the upper, AR-coated, opposite face of the beamsplitter plate.

Appendix J. The mathematics of not-quite-quadrature signals

The quadrature interferometry you have come to appreciate generates two non-redundant signals $S_1(t)$ and $S_2(t)$ from the standard and non-standard outputs of a Michelson interferometer. You have seen the value of an XY-plot of these signals on an oscilloscope for displaying the 'optical phase' of an interferometer in real time. Here's a mathematical treatment which explains the 'locus' you see on the 'scope, and how you can extract detailed quantitative information from it.

We'll think of a two-arm interferometer in which the 'optical lengths' of the two arms are L_1 and L_2, so that everything about the outputs depends on the interferometer's phase difference $\phi_{int} = k(L_1 - L_2)$. This 'optical phase' might vary due to the motion of either mirror, or for other reasons; for example, if there's a gas cell in one arm, then we've modeled the optical phase by

$$\phi_{int} = \phi_0 + \frac{2\pi}{\lambda}(n - 1)(2L) \quad ,$$

where n gives the refractive index of the gas in a cell of (one-way) length L placed in one arm of the interferometer. The important point about ϕ_{int} is that it's a variable, and can in some experiments be made to vary over a large multiple of 2π radians; in those cases, we'd 'watch the fringes go by' as ϕ_{int} varied systematically.

Now in the quadrature interferometer, the two output signals are both periodic in this phase ϕ_{int}, though they have different offsets and amplitudes, and (most importantly) have a phase difference between them. We could call that phase difference $\Delta\phi$, and we could allocate it half-and-half to each of the two distinct signals by writing

$$S_1(t) = a + A \sin\left(\phi_{int} + \frac{\Delta\phi}{2}\right) ,$$

$$S_2(t) = b + B \sin\left(\phi_{int} - \frac{\Delta\phi}{2}\right) .$$

Ideally the metal-film beamsplitter would give a perfect quadrature signal by fixing $\Delta\phi = \pi/2$, but in practice $\Delta\phi$ will be some angle near, but not right at, this value. The goal of this appendix is to show how that $\Delta\phi$-value can be extracted from experimental signals, and how (once it is known) the interferometer phase ϕ_{int} can *also* be extracted from those signals.

What you might have already done on the 'scope is to use the dc-offset controls to eliminate the effects of offsets *a* and *b*, to use scaling adjustments to make coefficients A and B equal, and then to use a 'scope's XY-plot to display the results in real time. That's equivalent to displaying the point

$$(x, y) = \left(\frac{S_1 - a}{A}, \frac{S_2 - b}{B}\right) = \left(\sin\left(\phi_{int} + \frac{\Delta\phi}{2}\right), \sin\left(\phi_{int} - \frac{\Delta\phi}{2}\right)\right) \quad ,$$

and watching its location change in the xy-plane as the interferometer phase ϕ_{int} varies. What is the locus of such xy-points?

The easiest way to find out is to compute x^2, y^2, and xy for this model, and to show that

$$x^2 = \sin^2 (\phi_{int} + \frac{\Delta\phi}{2}) = \frac{1}{2} (1 - \cos (2\phi_{int} + \Delta\phi)),$$
$$y^2 = \sin^2 (\phi_{int} - \frac{\Delta\phi}{2}) = \frac{1}{2} (1 - \cos (2\phi_{int} - \Delta\phi)),$$
$$x \cdot y = \sin (\phi_{int} + \frac{\Delta\phi}{2}) \sin (\phi_{int} - \frac{\Delta\phi}{2}) = \frac{1}{2} (\cos \Delta\phi - \cos 2\phi_{int}) \quad ,$$

and then to do the trigonometric manipulations which show that the quantity

$$x^2 - 2 x y \cos \Delta\phi + y^2 = 1 - \cos^2 \Delta\phi = \sin^2 \Delta\phi$$

is a *constant*. So as the interferometer phase ϕ_{int} varies, this combination of x- and y-values does *not* change. In fact, the constancy of this 'quadratic form' describes a set of xy-points forming an *elliptical* locus in the xy-plane. The ideal case of $\Delta\phi = \pi/2$ gives for the locus the familiar equation $x^2 + y^2 = 1$, just the unit circle, but the general case is a tilted ellipse in the xy-plane instead. The usefulness of this ellipse is that its shape will fix the phase difference $\Delta\phi$, while the xy-point's instantaneous location on the ellipse will give the interferometer phase ϕ_{int}.

Here's how to extract the phase difference $\Delta\phi$ from the raw signals $S_1(t)$ and $S_2(t)$ that are available from the two outputs of the interferometer. We suppose that we can digitally capture these waveforms while the interferometer phase is being varied linearly in time (for example, while the air pressure in the gas cell is rising linearly). On a digital 'scope, we might see several whole fringe cycles go by, and we might sample many points per cycle. Such data can be captured in a spreadsheet, and edited to length to cover some integer number of cycles (one, or many). Here's an algorithm for extracting information from such a spreadsheet:

- First, according to the model above, you can take the column averages (over any integer number of cycles) to give the values of the offsets *a* and *b*.
- Next, you can form new columns corresponding to $S_1 - a$ and $S_2 - b$.
- Third, you can form more new columns defined by $(S_1 - a)^2$ and $(S_2 - b)^2$, and these columns, according to the model above, will have average values $A^2/2$ and $B^2/2$. From these averages, you can extract the values of A and B.
- Now, with A- and B-values on hand, you can form columns defined by $C_1 = (S_1 - a)/A$ and $C_2 = (S_2 - b)/B$. According to the model above, these are described by

  $$C_1 = \sin (\phi_{int} + \frac{\Delta\phi}{2}), \;\; C_2 = \sin (\phi_{int} - \frac{\Delta\phi}{2})) \quad ,$$

- So you can form yet another column, defined by C_1 times C_2. It is giving you

  $$C_1 \, C_2 = \sin (\phi_{int} + \frac{\Delta\phi}{2}) \sin (\phi_{int} - \frac{\Delta\phi}{2}) = \frac{1}{2} (\cos \Delta\phi - \cos 2\phi_{int}) \;\; ,$$

  so *its* column average (taken, as all these averages are, over any integer number of fringe cycles) gives you $(\cos \Delta\phi)/2$, which finally gives you the value of $\cos \Delta\phi$ you need.

Note that with *no* least-squares fitting, you have extracted the values of all five parameters a, b, A, B, and $\cos \Delta\phi$ that describe the model; note too you have incorporated *all* the data points, not just special points like the peaks and valleys.

With all these parameters extracted from the xy-locus of the fringe signals, it is now possible to return to the instantaneous values of the recorded signals and extract, point by point, the time-varying value of the interferometer phase ϕ_{int}. The easiest method involves yet more columns in your spreadsheet, since you can show that

$$\frac{1}{\sin \Delta\phi}(C_1^2 - C_2^2) = \sin 2\phi_{int} \quad ,$$

$$\cos \Delta\phi - 2\, C_1\, C_2 = \cos 2\phi_{int} \quad .$$

The left-hand side of these equations can be computed, row-by-row in the spreadsheet. [You'll note that you're in trouble if that denominator, namely sin $\Delta\phi$ is too small; but ideally it'll be one, and in practice it'll be far enough from zero.] The right-hand side of these equations gives what you need to compute tan $2\phi_{int}$ on a *point-by-point* basis, so you have finally gotten a way to extract the interferometer phase separated from all other instrumental effects.

Now if you plot the values of ϕ_{int} you have obtained as a function of time, you will see the 'phase unwrapping' problem confront you. The difficulty is that the inverse-tangent function has a range of only π radians (in $2\phi_{int}$), so your value of ϕ_{int} is off by some multiple of $\pi/2$ radians. It's very easy to see by eye, by continuity in a graph, what multiple of $\pi/2$ needs to be added to any given part of the data, but it's harder to automate that process. The problem is that your line-by-line computation of the interferometer phase fails to take into account what quadrant the point (cos $2\phi_{int}$, sin $2\phi_{int}$) actually lies in, and it also fails to count the number of times the moving point (cos $2\phi_{int}$, sin $2\phi_{int}$) has wound itself around the origin as the interferometer phase changes. You might want to 'unwrap the phase' by hands-on correction of entries in your spreadsheet, or think about how you can compute sample-to-sample phase changes to follow the phase continuously; or you might look into how the 'phase unwrapping' problem has been handled in the research literature.

Appendix K. Hysteresis -- its usefulness and detection

The goal of this section is to explain some of the subtleties that are involved in counting digitally the successive fringes that might appear in a Michelson interferometer when it's being scanned. The difficulties, and their solutions, are also more broadly applicable to any situation in which an analog signal is being cycle-counted by digital electronics.

We suppose that a Michelson interferometer is being motor-driven to produce about 2 fringes per second, so that an analog output waveform is sinusoidal and roughly periodic. We further suppose that it's had a hand-adjusted zero-offset correction applied, so that the signal is roughly symmetric above and below zero. Now one goal would be to count the number of full cycles that go by in a specified time, or between two specified endpoints of travel read off a micrometer. If there are hundreds or thousands of counts, it's tedious and potentially inaccurate to do this 'count' by eye-hand-mind coordination, so we might resort to digital means. What we might use, on that analog signal, is a 'comparator' circuit, whose response depends only on the *sign* of the voltage at its input. Applied to the waveform discussed, the comparator's output would be

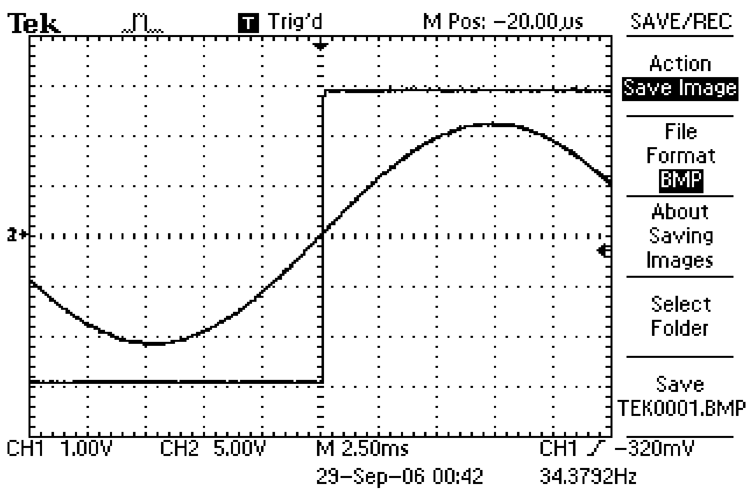

Figure K-1: Views of input, and output, waveforms for a comparator

Now the comparator's output signal has sharp well-defined transitions, and one could imagine a digital counter which incremented its count by +1 for every rising edge of the comparator's output. That would give a desired +1 count per original cycle of analog input.

The reason that this would *not* work reliably in practice has to do with timescales and with electronic noise. Recall that we're assuming analog sinusoids with frequency near 2 cycles per second, or periods of about 500 ms per full cycle; as a result, the analog waveform spends a time of order *milliseconds* in the vicinity of each zero-crossing. The comparator circuit might have a much shorter time response, of order tens of *nanoseconds*. That mismatch by itself would not present a problem, except that real analog signals are always more or less noisy; in this case, both 'real' optical-caused fluctuations and mere electronic noise make the signal, in the vicinity of a zero-crossing, look more like this:

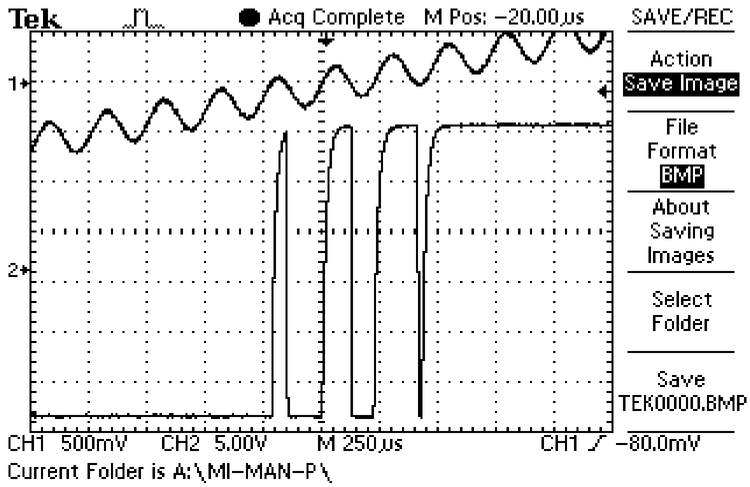

Figure K-2: Views of input, and output, waveforms for a comparator, in the presence of a waveform showing a slow up-slope and sinusoidal fluctuations, simulating a view of an optical fringe slowly crossing up through zero in the presence of additional 'noise'.

It should be clear that there can be *many* actual negative-to-positive zero-crossings occurring during one nominal rise of the analog signal up through zero. Many more such negative-to-positive zero-crossings also occur during the nominal *fall* of the analog signal downwards through zero, half a cycle later. So one actual analog cycle will result in an unpredictably large number of digital counts.

One tool that can be used to deal with this noise is a low-pass filter, which suppresses the high-frequency content in the analog waveform and thereby removes some of the wilder fluctuations that result in multiple counting. But analog filtering also induces a time delay, sometimes undesirable in itself, and in any case it doesn't get to the heart of problem. What's used instead is a comparator with hysteresis, whose output state depends not only on the sign of the input voltage, but *also* on its past history.

That behavior is best shown in a diagram of output state as a function of input voltage for two comparator circuits. On the left is the simple comparator, for which this diagram has a sgn(x) sort of discontinuous behavior. On the right, is the curve for a comparator with hysteresis, where the output depends on the *direction* of travel of the input voltage, here depicted by an arrow indicating sense of change with time.

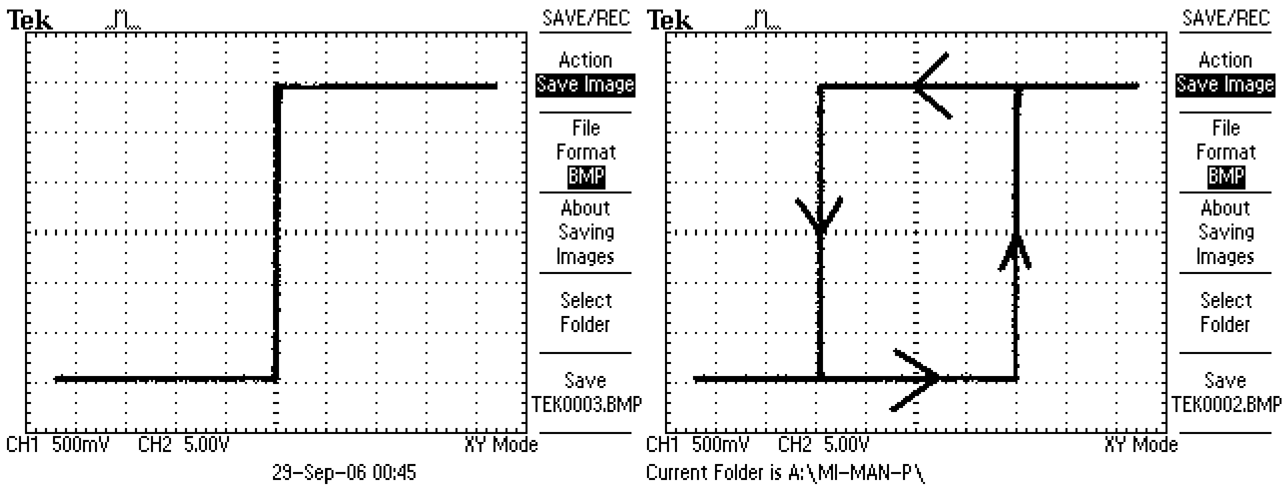

Figure K-3: X-Y oscilloscope plots with input on X-axis (Ch1) and output on Y-axis (Ch2): on left hysteresis is off, and on right hysteresis is 1.0 volts

So now it's clear that a negative-but-increasing analog voltage will cause a digital transition, not at the instant the input voltage reaches zero, but later when it reaches +Δ. What's more, once that upward-doing digital transition has occurred, no second transition can occur until and unless the input voltage has first reached as low as -Δ. The size of Δ, which measures the hysteresis, can be chosen to exceed the plausible range of the analog noise -- though of course Δ has to remain smaller than the amplitude of the input analog sinusoid for counting to work at all.

It is well worth seeing this process working with actual interferometric fringes. So set up a Michelson interferometer with motor drive in action, such that you can get an ongoing set of steady fringes. Put those fringes into an input on your controller box, and use an oscilloscope to view the fringes. First monitor the effect of the zero-offset control to ensure you have a sinusoid more or less centered about zero, and then monitor the effects of a low-pass filter on the signal. If you pick a 'time constant' too large, you will markedly reduce the size of your sinusoid, so go to a smaller, or even minimal, time constant. Now use a BNC cable send your signal onward to the channel-A comparator in your controller, and arrange your 'scope to view, in dual-trace mode, both the input to, and the output from, this comparator. Here are two ways to see that the comparator is working:

- You can arrange the 'scope to trigger on (say) the upward-going transition of the comparator *output*, and then also view the state of the analog *input* which makes this occur. [You expect to see the analog signal increasing through the value +Δ (not zero) at the time of this transition.] For another look, check what happens if you have the 'scope trigger on the *downward*-going transition of the comparator output.
- You can create an untriggered continuous XY-display on your 'scope, with the analog input voltage traversing the X-axis, and the comparator's output voltage displayed on the Y-axis. You should get a direct view of the hysteresis diagram, and can very directly see the effect of setting different Δ-values for your comparator. You can also see what happens if you pick a value of Δ too large given your signal's amplitude, or if your zero-offset is incorrectly adjusted.

By either of these methods, you should be able to adjust the hysteresis to give reliable counting on the digital display on your controller, by using the 'channel A' (rather than 'quadrature') mode of counting. In particular, you should confirm, by cross-comparison of the analog input signal on the 'scope and the state of the counter, that the counter increments once (not twice, and not many times) per cycle. Similarly, when you come to count the signals from a quadrature Michelson interferometer, you can use hysteresis (rather than low-pass filtering) to ensure that you get only one count (up or down) when the analog input state crosses quadrant boundaries.

Appendix L. Choices for table-mounting the mirror mounts

The TeachSpin mirror mounts you've been using all involve an 'upright' mounted on a 'baseplate' of 1" thickness, and the orientation of the mirror is controlled by the way you mount the baseplate to an optical breadboard. Here's a discussion of why there are no fewer than *nine* holes through the baseplate, and how they can be used.

Near the center of the baseplate, vertically under the center of the mirror, is a single hole that allows a 1/4-20 socket-head cap screw to hold the baseplate down to the breadboard. If you use this single-screw mount, the mount can be rotated continuously, allowing the mirror to reflect an incident laser beam through any desired angle.

But more often, you'd want a particular choice of angle between the beams incident on, and reflected from, the mirror. In Michelson interferometers, you want retroreflection; that requires a mirror to be in the 'face-on' configuration. In Sagnac interferometers, you want the beams deflected through a 90° angle; that requires the mirrors to be turned by 45° relative to the face-on configuration.

Those particular orientations, and one other, can be accommodated by using 1/4-20 screws in two particular holes in the baseplate. The figures below show the technique: note the holes that are used will align with the 1"-square grid work of holes on an optical breadboard.

- for 'face-on' geometry and direct retroreflection --

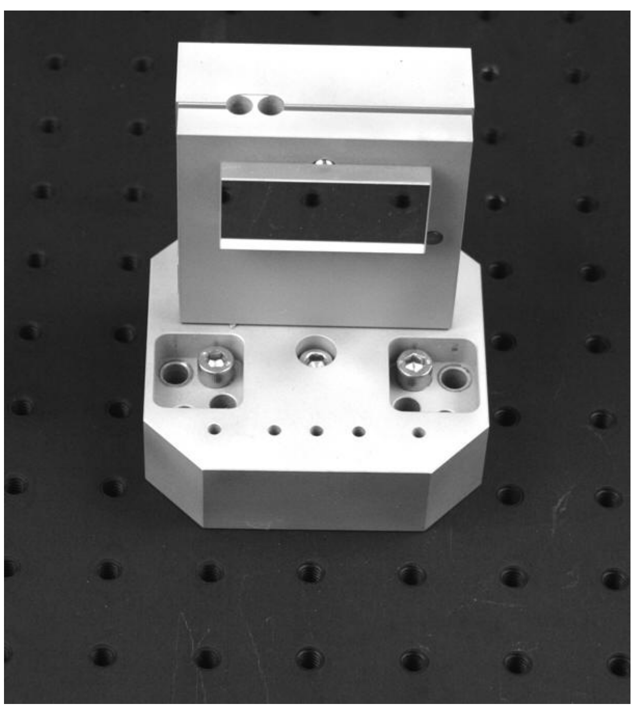

Figure L-1: Two bolts at 2.00" spacing hold mirror mount for retroreflection

- for a mirror rotated by 45°, and a beam deflection through 90°--

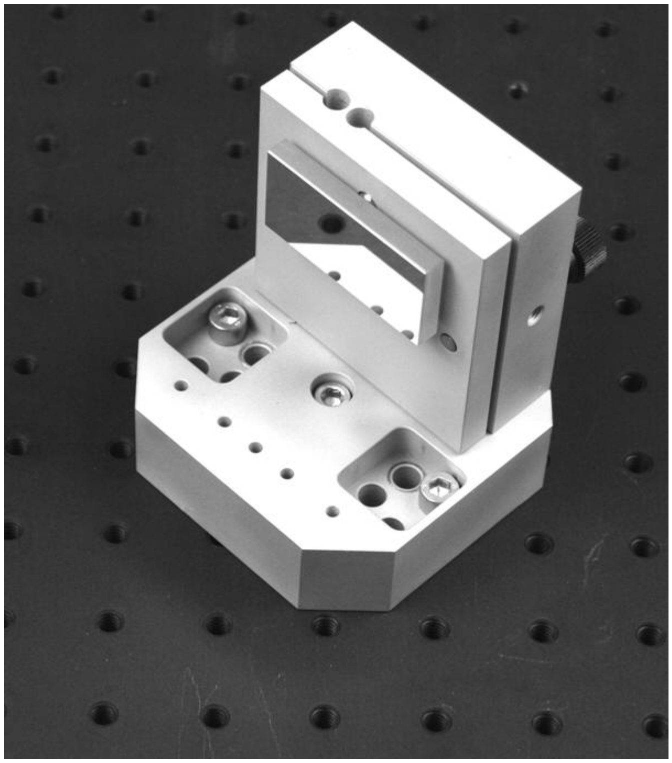

Figure L-2: Two bolts at 2.828" spacing hold mirror for 90° deflection

- for a mirror rotated by 22.5°, and a beam deflection through 45°--

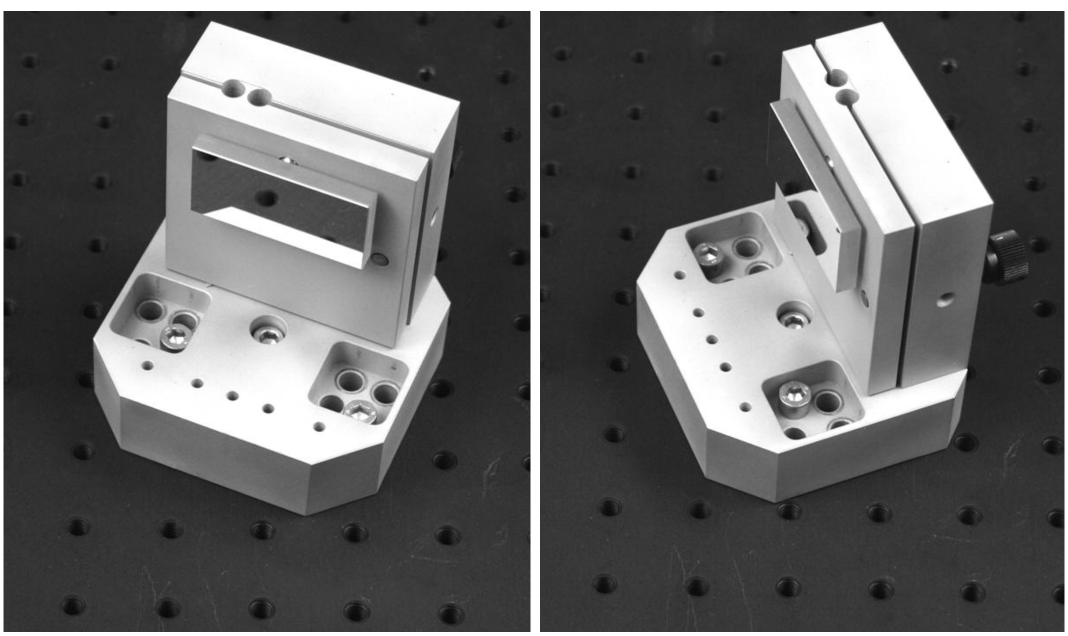

Figure L-3: Two bolts in base hold mirror mount for 45^{O} deflection, two possible orientations shown

Note that in all these cases, the single central hole is still located above a hole in the optical breadboard, so in all these cases the mirror has been rotated about a single central axis. Note that when *pairs* of screws are used as in the configurations above, the central hole may be left empty of a screw. Note also the tremendous utility of a ball-tipped Allen driver for the socket-head cap screws; without such a tool, it would be all too easy to damage the mirror itself while wrenching around with the mounting screws.

Finally, the TeachSpin mirror mounts are driven, in their single degree of rotational freedom, by an 80-pitch screw. The line of action of that screw is about 1.74" from the hinge line of the flexure, so that rotating the adjusting screw by one full turn gives an angular deflection of the mirror by angle (1/80)"/1.74" = 19.7 mrad. The angular deflection of a beam of light reflecting off the mirror is double this, or about 39.5 mrad per full turn. Of course, adjustments of a tiny fraction of a full turn are feasible.

Appendix M. Straight-line fringes and spatial frequency

Particularly in the Sagnac interferometer, you will have noticed that the output beam of the interferometer, suitably viewed on a screen, appears to be crossed with a number of parallel dark stripes, or dark fringes. Why are there these 'fringes'? and why are they straight? and what does their spacing depend upon? This appendix models such fringes.

We start by assuming that the two beams emerging from the Sagnac interferometer can both be treated as plane waves. (In practice, their wavefronts are spherical, and have small and equal curvature; and in practice, the wavefronts are transversely truncated, to form Gaussian spots of several mm in diameter.) Clearly the beams need to overlap in position to give the destructive interference that causes the dark fringes; we will see that it's imperfect overlap in *direction* of propagation that controls the fringe spacing.

To see this in the simplest geometry possible, consider two waves of the same frequency which are moving approximately in the +x-direction, but which are separated by angle θ in their directions of propagation. The wave vectors of the two waves both have magnitude

$$k = \frac{\omega}{c} = \frac{2\pi f}{c} = \frac{2\pi}{c/f} = \frac{2\pi}{\lambda} \ ,$$

but they can be taken to have directions making angles of ± θ /2 relative to the x-axis:

$$\mathbf{k}_+ = k(\hat{x} \cos\frac{\theta}{2} + \hat{y} \sin\frac{\theta}{2}) \ ; \ \mathbf{k}_- = k(\hat{x} \cos\frac{\theta}{2} - \hat{y} \sin\frac{\theta}{2}) \ .$$

Here the misalignment is taken to be along the y-axis. Now if the two waves have equal amplitude and the same polarization, then the complex representation of the overlapped waves is

$$\mathbf{E}_{sum}(\mathbf{r}) = \mathbf{E}_0 \exp(i\, \mathbf{k}_+ \cdot \mathbf{r} - i\, \omega\, t) + \mathbf{E}_0 \exp(i\, \mathbf{k}_- \cdot \mathbf{r} - i\, \omega\, t) \ .$$

Using the representation above for the k-vectors, this can be written as

$$E_{sum}(x, y, z) = E_0 \exp(i\, k\, x \cos\frac{\theta}{2}) \left(e^{i\, k\, y \sin\, \theta/2} + e^{-i\, k\, y \sin\, \theta/2}\right) e^{-i\, \omega\, t}$$

$$= E_0 \exp(i\, k\, x \cos\frac{\theta}{2})\ 2 \cos(k\, y \sin\frac{\theta}{2})\, e^{-i\, \omega\, t} \ .$$

The intensity of the wave is proportional to the absolute square of this function, which gives a predicted intensity

$$I\,(x, y, z) = 4\, E_0^2\, \cos^2(k\, y \sin\frac{\theta}{2}) \ .$$

For the waves in question, a viewing screen would be placed in a plane with x = constant, and with y- and z-axes lying in the screen plane; so this function describes illumination varying (only) in the y-coordinate across the screen. In fact it predicts bright fringes whenever the cosine reaches value ±1, i.e., when the argument of the cosine function is nπ, for any integer n. So the bright fringes are located at

$$k\, y_n \sin\frac{\theta}{2} = n\, \pi \ ; \text{ or } y_n = \frac{n\, \pi}{\frac{2\pi}{\lambda} \sin\frac{\theta}{2}} = n \frac{\lambda/2}{\sin\theta/2} \ .$$

So the fringes are predicted to be straight lines, and equally spaced, with fringe spacing

$$\Delta = \frac{\lambda/2}{\sin \theta/2} \quad .$$

In practice, the misalignment angle θ will be of order of a milliradian (if the laser beams used are going to overlap in position at all) so the small-angle approximation gives $\Delta \approx \tilde{\lambda}\tilde{\theta}$

If the angular misalignment of the two beams is a full milliradian, then we get $\Delta = 10^3\ \lambda \approx 0.6$ mm, typical of the fringe spacing seen during initial alignment of an interferometer. As the beams are adjusted to come more nearly into angular alignment, the angle θ decreases and the fringe spacing Δ increases.

It is often useful to think of Δ as giving 'spacing per unit fringe', and to introduce its reciprocal $1/\Delta$ as giving 'fringes per unit distance', or *spatial frequency* of fringes. Since $1/\Delta \approx \theta/\lambda$, that quantity is directly proportional to the angular misalignment, so as the beams are brought into perfect angular alignment, and the misalignment angle θ passes through zero, the spatial frequency of fringes also goes through zero. It's easy to see the spatial frequency of fringes start high, then zoom toward and through zero, and then head back up to large values, as some adjustment knob is turned. For the more realistic case in which **k**-vectors are misaligned in both y- and z-senses, it is easy to use one knob to adjust (say) the y-misalignment to give minimal spatial frequency, and then to use the orthogonal knob for the z-misalignment to drive the spatial frequency all the way to zero.

Appendix N. The mounts for polarizing beamsplitter cubes

Your first Sagnac interferometer will use a polarizing beamsplitter cube (PBSC) held in the incident beam at normal incidence. This is accomplished using a beamsplitter-cube holder; it recycles the same octagonal base as the holder for 1"-diameter thin-plate beamsplitters, but has a different upright attached to it.

[Appendix H treats the functions and use of the thin-plate beamsplitters.]

The polarimetric detection scheme for the Sagnac interferometer uses a *second* PBSC held at a curious angle in a special jig, and you will need to transfer one of your two cubes from a beamsplitter-upright into this special jig. This appendix also contains directions on how to accomplish that task safely.

The base-and-upright combination is designed to hold a 1" PBSC at the proper height and orientation above the optical breadboard. In the upright, the glass cube is held between a metal frame (below) and a black plastic cover (above). Only if you have an upright in its working vertical orientation is it safe to remove the cover -- you don't want to loosen the cover otherwise, lest the glass cube fall out and land on one of its vulnerable corners.

If you have an upright oriented safely, you may use a 5/64" ball-tipped Allen driver to remove the four Allen-drive flathead screws retaining the black plastic cover. When you lift the cover off vertically, you'll see the top face of the PBSC, and the internal hypotenuse that bears the active beamsplitting surface of the device. [Note that the orientation of that diagonal hypotenuse is as shown in the diagram below -- not rotated through 90° as it might be. This is to make the correct *location* of the cube more obvious.]

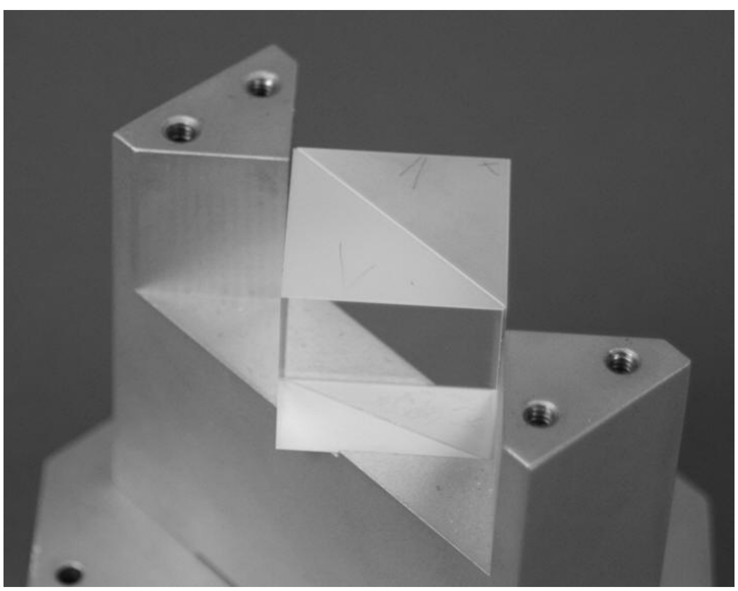

Figure N-1: The proper orientation for a PSBC on its beam splitter mount (with the black plastic cap removed)

The cube is now in a position to be removed, **but first please understand how to handle a PBSC properly**. It's a glass cube with four polished optical faces, and (as it stands now) four vertical edges. You are NEVER to handle or touch any of these faces or edges, whether with fingertips or gloved hands; rather, you are to engage the cube only by its frosted-glass top and bottom surfaces, where fingerprints and scratches can do no harm. Similarly you're only to set the cube down with a frosted surface bottommost.

If you heed this warning, you may lift out the cube, and you'll see the shallow diagonal groove in the metal upright that helps to orient the cube properly. Now if you look at the underside of the plastic cover, you'll see another such shallow diagonal groove. By design, these two grooves end up being oriented perpendicularly, so that the glass cube will be held in such a way that it can't fall or slide out.

Now practice placing the cube in the groove in the upright, orienting it so that its internal hypotenuse is lined up correctly between the two metal edges that nearly touch its ends. Then add the plastic cover, ensuring that its diagonal groove also captures the cube. When that's in place, you can engage all four flathead screws to hold it in place. When tightening those screws, the goal is **NOT** to close the gap between plastic cover and metal frame, but rather to put a *gentle* bend in the plastic cover to ensure the cube is held firmly but safely.

Now that you know how to remove and re-install a PBSC safely into the beamsplitter uprights, here's how to transfer a PBSC into the special polarimetric jig you'll use in the detection of signals from the Sagnac interferometer. That jig too holds a PBSC between a metal base and a black plastic cover, and in that jig too there are shallow grooves in the base and the cover to engage the glass cube. The jig is shipped to you without a cube installed, so practice removing its cover. You'll note that one of the four screws is different from the other three; that's intentional, and designed to ensure that (again) the two shallow grooves for the cube end up perpendicular to each other.

When you've learned how the cube-holder works, and when you have safely removed a PBSC from a beamsplitter upright, you are ready to install it into the Sagnac-detector jig. Again remember to handle the cube always and only by its frosted-glass top and bottom faces, never by its polished faces. When you install the cube, it's very desirable to have the jig tilted, as shown below, so that you have a level base to set the cube upon.

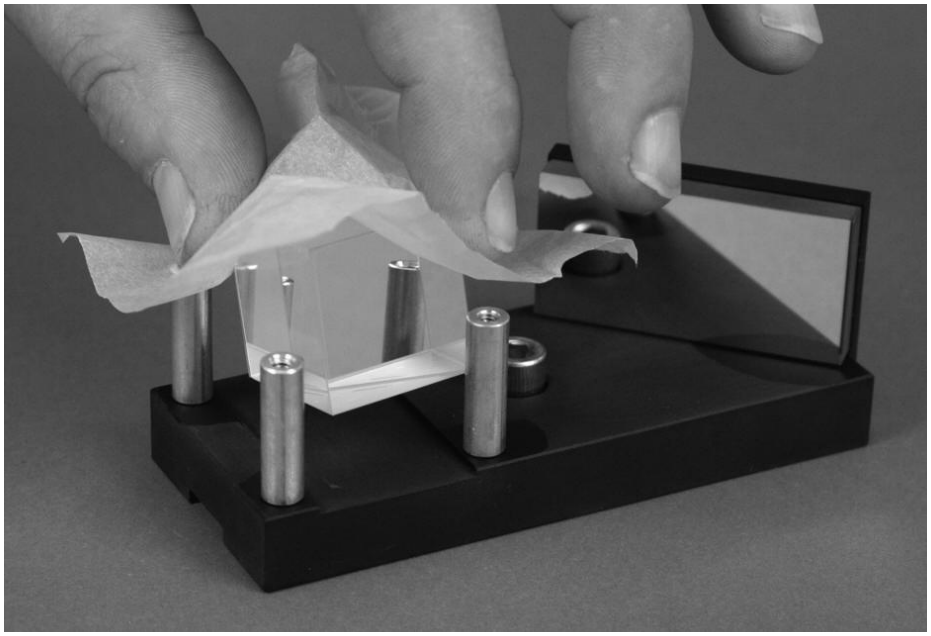

Figure N-2: Installing a PBSC into the Sagnac-polarimetric arrangement

And when you install the cube, make sure its internal hypotenuse ends up in the orientation shown (parallel to the fixed front-surface mirror near it) so that it directs light in a useful direction.

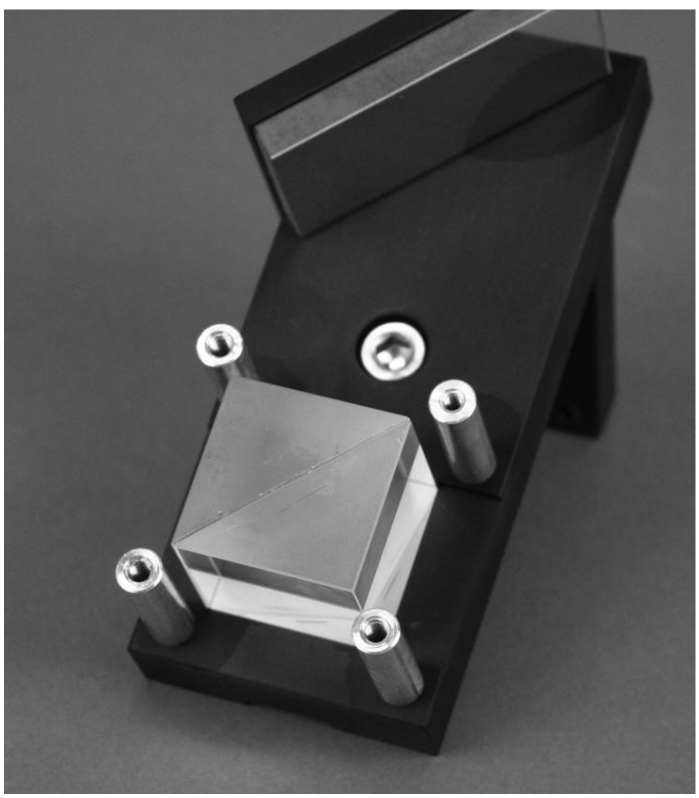

Figure N-3: The proper orientation of the internal beam-splitting hypotenuse in a PBSC

With the cube in position, you are ready to put on the top plastic cover; ensure that its groove is perpendicular to the groove in the base you've just used. Then engage the four screws that hold the cover in place, and tighten them until you feel them snug the plastic against the PBSC -- you don't need to use them to crush or distort the glass cube.

Appendix P. The miniature rotation stage

Your TeachSpin interferometry kit includes one miniature 'rotation stage', which permits you to give carefully-controlled rotations to a vertically-oriented 1/2"-diameter optical post (and anything attached to it). You might use this capability to rotate slabs of glass or quartz for index-of-refraction studies, or the compensator plate for white-light interferometry.

The rotation stage as supplied has an OptoSigma rotator unit sandwiched between a standard base and a post-holder, so it can be clamped down to the optical breadboard, and carry an optical post, as usual. A black knurled thumbscrew is the way to clamp onto a 1/2" mounting post, and also to adjust the height of the optical component that it bears.

The rotator proper has three silver knurled knobs, the thinnest of which can be loosened to unlock the post-holder for free rotation about a vertical axis. Once it's loosened, the post-holder can be rotated freely through any number of full turns; meanwhile, you can read its angular position on a 0-360° angular scale.

When this knurled knob is tightened, the upper post-holder is locked to the lower rotator unit. Thereafter the post-holder can be rotated incrementally, by about ±5°, from the previously chosen position. The way to get that incremental rotation is to turn the knurled knob that is connected to a threaded rod passing through a bronze sleeve. Viewing the rotator from the side, you can see that threaded rod's end bearing on a sort of 'lever arm' and thereby turning the whole upper part of the device.

[The third knurled silver knob is the tension adjustment for the push-spring that keeps the lever arm engaged against the threaded rod.]

The threaded rod has a thread of pitch 0.50 mm, and it bears on the lever arm at a distance of 17.5 mm from the axis of rotation. So (at least in the vicinity of perpendicularity) the knob-and-rod rotates the arm-and-post holder by (0.50/17.5) = (1/35)th of a radian per turn of the knurled knob. Subdivisions of full turns will give even finer rotational adjustments.

There's also a black plastic rod with a black plastic sleeve attached to one end; this 'wand' may be attached to the knurled knob simply by press-fitting the sleeve over the knob. The purpose of the wand is to allow access to the rotation stage's angular adjustment even when the rotator is part of an interferometer and the draft shield is covering the rotator and its contents. To use the wand in this situation, you'll need to position the rotator in such a way that the wand can exit the draft shield via one of the little 'mouse holes' you'll find around its lower edge.

Appendix Q. Preparing a thermal-expansion sample

Your TeachSpin equipment includes one fully prepared thermal-expansion sample, and the parts required for you to prepare more. In fact, you can prepare a sample from any material which you can shape into a rod of about 1/4" or 6 mm diameter, at least 3" or 75 mm length, with flat and plane-parallel ends. The only requirements are that

- the sample have adequate strength in compression to serve in the mechanical train illustrated in Figure 12-1, and
- the sample have adequate thermal conductivity, so that the temperature measured on its surface should apply adequately to its full cross section.

Here's how your copper sample was prepared for your use; the ingredients are

- the sample rod itself
- a spool of teflon-insulated constantan heater wire of 0.010" diameter
- connecting wire for the heater, in the form of a flat-ribbon pair
- a type-K (chromel-alumel) thermocouple, already ‘connectorized’
- a supply of 3/8"-wide Kapton tape
- heat-shrink tubing

The procedure is as follows:

1) Cut a 2-foot, or 60-cm, length of heater wire, and wire-strip its ends.
2) Wire-strip both wires in the ribbon pair, at both ends, as shown below in Figure Q-1. Note that you want the stripped ends ‘offset’ on one end of the ribbon pair.

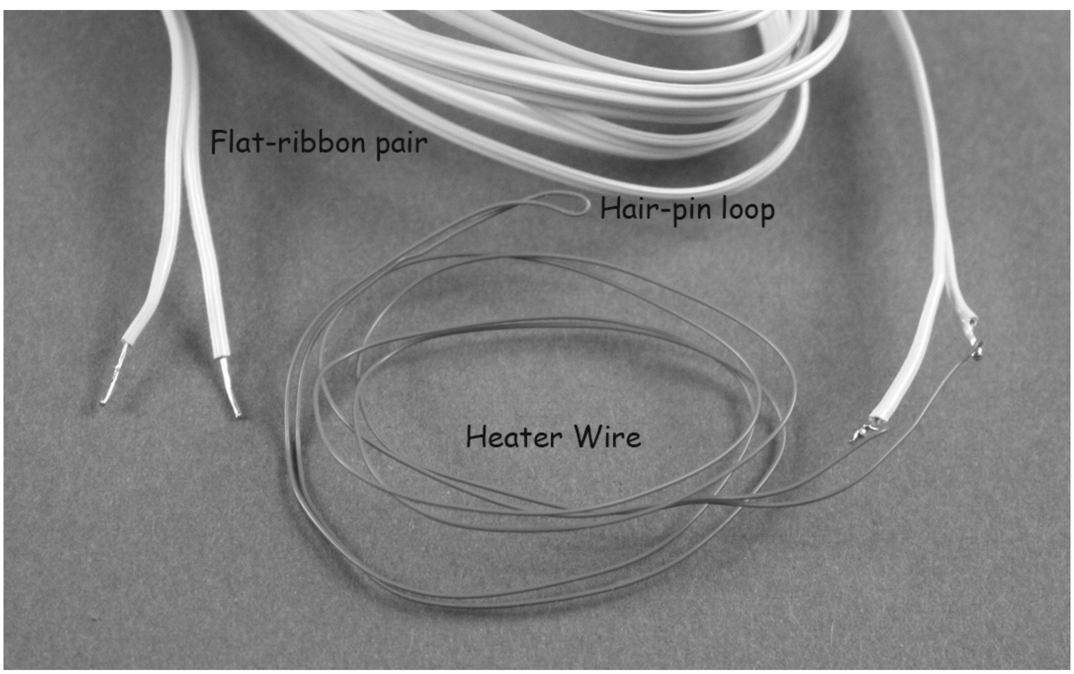

Figure Q-1: The heater and flat-ribbon pair wires properly terminated.

3) At the ‘offset’ end, twist-connect each of the heater wire’s stripped ends to the wires of the ribbon pair.
4) Now solder the twist-connections you’ve just made and also ‘tin’ the exposed wires at the opposite end of the flat-ribbon pair.
5) Pull out a 'hairpin loop' in heater wire, forming a pair of parallel wires, each about 30 cm long. You will be winding this loop of paired wires around the sample rod. The idea of the hairpin is to let the 'start' and 'finish' ends of the heater to come out at the same end of the sample rod, and also to minimize any magnetic effect of the heater current.
6) You want the hairpin loop of heater wire to form a ‘double helix’ around the sample rod, covering nearly all of its length. Prepare a spot near one end of the sample rod by laying down about 1-1/2 turn of Kapton tape pad; this will prevent the heater wire’s soldered connections from making contact with the rod (or with each other).
7) On the opposite end of the rod, tape the hairpin loop end of the heater wire to the rod using about 1-1/2 turns of Kapton tape. Continue winding the wire the length of the rod keeping the spacing between turn-pairs as uniform and as taut as possible. It is important that the heater wires in the pair do not cross each other as this could lead to a possible failure point in the heater wire.
8) Ensure that the soldered junctions of the heater wire/flat-ribbon pair wires lie on top of the Kapton tape pad, applied in step 6 above. Take another piece of Kapton tape about 2-inches long and carefully tape the connections to the rod. Make sure they do not touch each other or come into contact with the rod. Make a 90-degree bend in the flat-ribbon pair close to the last solder joint, facing the center of the rod and tape it securely in place as shown in Figure Q-2.
9) Near the center of the sample, spread the heater wires sufficiently to make room for mounting the active tip of the thermocouple. Tape it into place right onto the sample rod with a piece of Kapton tape about 1" or 25mm long. It is *NOT* necessary to isolate the thermocouple’s tip electrically from the sample. The thermocouple should be positioned along the length of the rod with the main part of the wire towards the heater connection as shown in Figure Q-2.

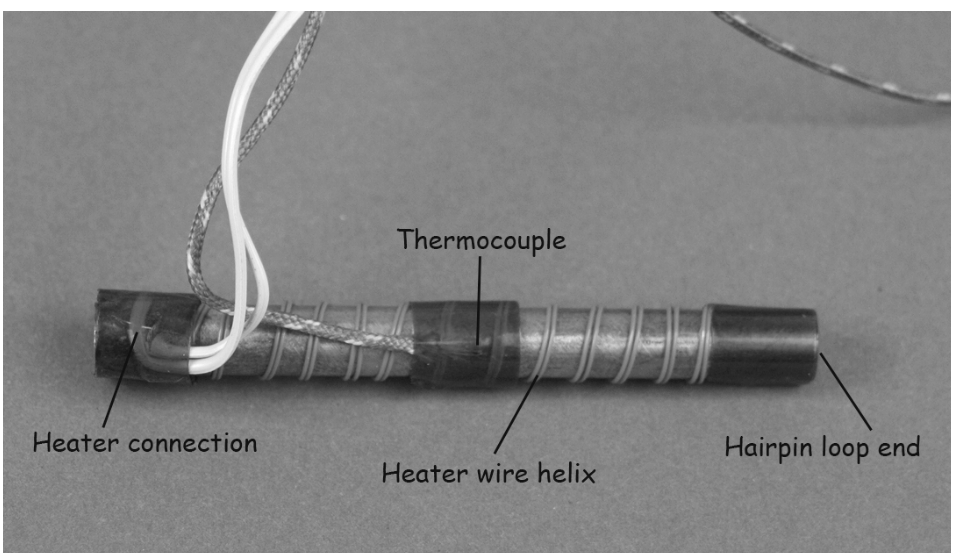

Figure Q-2: Sample rod with heater and thermocouple installed.

10) Test your heater for electrical continuity, about 6 ohms, and for electrical isolation from the sample rod. Also test that your thermocouple works, and will read out a believable temperature.
11) Slide a 2-1/2" or 60 mm length of heat-shrink tubing over the hairpin loop end of your sample, and use a heat-gun to shrink it down over the heater (and thermocouple connections) flush with the rod end. Bend the thermocouple wire perpendicular to the rod and orient the flat-ribbon pair towards the thermocouple wire and bend it perpendicular to the rod at the thermocouple wire. Take another piece about 3/8" or 10 mm long and shrink this down over the connection end of the rod sample and shrink in place. Refer to Figure Q-3 for this detail.
12) When finished, carefully trim any excess heat shrink and tape away from the faces of the rod ends.

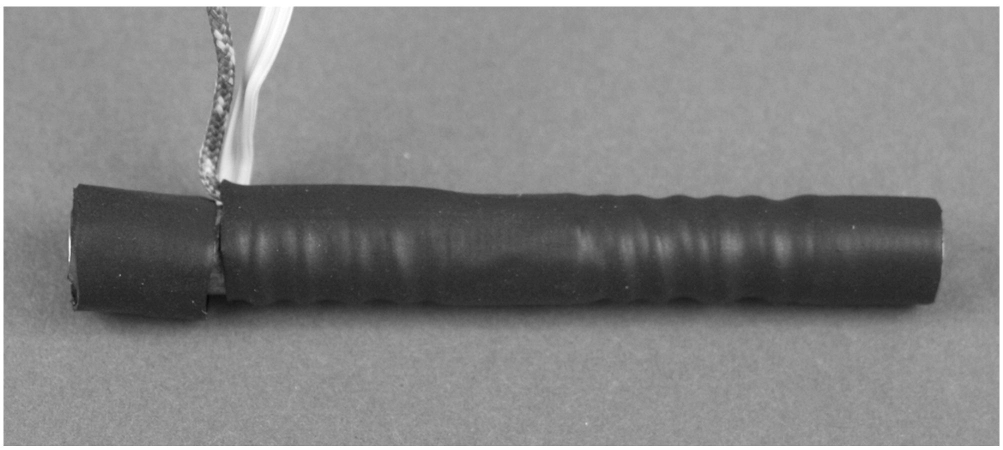

Figure Q-3: Finished sample rod assembly.

Your sample is now ready for use. Clearly you can undo any part of this procedure for repairs or revisions; perhaps you can also find ways to improve your technique. You might also think of trying to embed the tip of the thermocouple into a hole drilled into the sample rod; or you might add thermal insulation to the exterior surface of the finished sample.

Appendix R. The index of refraction of optical glass

'Glass' is not a well-defined chemical compound of definite proportions, but rather a whole class of materials formed from the oxides of silicon, calcium, sodium, and other elements. Glass that's optimized for transparency and homogeneity might be called 'optical glass', and suppliers of optical glass have many standardized recipes for the production of glasses of well-defined properties.

You might have heard of 'soda-lime' glass, with a high proportion of sodium and calcium oxides, or 'borosilicate glass' with boron and silicon oxides predominating. Or you may have heard of optical 'crown' and 'flint' glasses, each in turn describing a whole class of recipes. But for optical purposes, tighter specification of a glass type is required.

One classification is based on the lens-designer's central concern for refractive index, and depends on giving the index of refraction at three standard wavelengths:

$n_C = n(\lambda = 656.3$ nm, the hydrogen red Balmer-α or Fraunhofer 'C line')
$n_D = n(\lambda = 589.3$ nm, the midpoint of the yellow sodium 'D lines')
$n_F = n(\lambda = 486.1$ nm, the hydrogen turquoise Balmer-β or Fraunhofer 'F line')

Roughly speaking, n_D is a measure of the refractive index of the glass for visible light, and the difference $n_F - n_C$ is a measure of the *dispersion* of the glass. The dispersion is typically quantified by the Abbe v-value, where

$$v = \frac{n_D - 1}{n_C - n_F}$$

so that the Abbe v-value is larger for glasses of smaller dispersion.

Suppliers of optical glass can provide samples covering a subset of the plane described by (n_D, v) pairs. Such glasses also differ markedly in their chemical and mechanical properties, and their cost and ease of polishing. By far the most commonly used glass for generic optically transmissive components is designated by its Schott recipe number BK-7, a borosilicate crown glass with a whole range of desirable properties. Its indices of refraction are given approximately by

$$n_C = 1.51432,\ n_D = 1.51673,\ n_F = 1.\ 52238$$

so that its dispersion has $v = 64.2$. Another designation for this glass is formed from three digits each of $n_D - 1$ and v, so it would be labeled 517 642.

Glass produced according to the BK-7 recipe has to pass very stringent standards for predictability of refractive index, so it is possible for the manufacturer to publish a table of (six-digit!) refractive-index values as a function of wavelength. Such tables, in turn, can be summarized by a set of fitting parameters. The fitting function most closely tied to a physical model for refraction is given by a Sellmeier formula, in which $n(\lambda)$ is specified by

$$n^2(\lambda) - 1 = \sum_{i=1}^{3} \frac{c_i \lambda^2}{\lambda^2 - \lambda_i^{\,2}}$$

Over a more limited range in wavelength, the same table of data can be adequately represented by a more empirical formula for the *square* of the refractive index,

$$n^2(\lambda) = A_0 + A_1 \lambda^2 + A_2 \lambda^{-2} + A_3 \lambda^{-4} + A_4 \lambda^{-6} + A_5 \lambda^{-8}$$

Expressing the n-value of the glass relative to standard air, and expressing standard-air wavelengths in units of μm, this six-parameter fit to the index of refraction of BK-7 glass is specified by constants

$A_0 = 2.2718929$
$A_1 = -1.0108077 \times 10^{-2}$
$A_2 = 1.0592509 \times 10^{-2}$
$A_3 = 2.0816965 \times 10^{-4}$
$A_4 = -7.6472538 \times 10^{-6}$
$A_5 = 4.9240991 \times 10^{-7}$

This formula gives $n_D = n(\lambda = 0.5893\ \mu m) = 1.51673$, and in fact it correctly represents the index of refraction for BK-7 glass to a precision better than 10^{-5} over the full range of wavelengths 0.4 - 1.2 μm.

Appendix S. The gas-handling manifold, and uses of the pressure transducer

This appendix describes how to use the gas-handling manifold (the curious wooden board with brass fittings and four valves) together with the TeachSpin gas cell, and pressure transducer, to produce controllable and quantifiable pressure changes inside the gas cell.

Please note first of all a **SAFETY WARNING**: the gas cell can withstand an internal pressure anywhere below 1 standard atmosphere, all the way down to the vacuum condition of zero pressure. But it is NOT intended to withstand an internal (absolute) pressure much above 1 atmosphere. So (independent of what pressure difference you choose to measure) keep the pressure inside the cell in the range 0 to 1.1 atmosphere (absolute), or in the range -1.0 to +0.1 atmosphere (relative to ambient air). Motivation: you *don't* want internal overpressure to blow the glass end windows off the cell.

To use the gas cell, it is essential to have a pump which can provide some semblance of vacuum; the ordinary laboratory forepump, capable of lowering air pressure from 1 atmosphere = 10^5 Pascals = 760 Torr to a limiting pressure of under 50 mTorr, is the natural device to use. You can connect the pump's low-pressure (vacuum) port to the nipple-fitting adapter on the gas manifold via a suitable hose. You may want to mount the pump well away from your interferometer, and to mount your gas manifold off of your interferometer-bearing lab table, to minimize the coupling of vacuum-pump vibrations into your interferometer.

Before turning on your pump, you might want to close all the valves on your manifold. Each knob is attached to a right-handed screw thread, driving a shaft inward to close a valve, so it's fully *clockwise* rotation of a knob that *shuts* a valve. Finger-tight is tight enough; use of wrenches would likely damage the valve's internal O-ring.

Now there are three ways to think about what's going on in your valves and hoses.

- You can think of vacuum as some 'commodity', produced in your vacuum pump, and 'conveyed' along through tubing and opened valves to where you need to use it.
- You can use the concept of 'suction', except that upon reflection you'll understand that you can't pull on a gas!
- Or, you can realize that gas will expand to fill any volume accessible to it, including a volume newly accessible upon opening a valve, or the volumes that open up during the intake stroke of your pump. In this view, you are following the movement of molecules, not mythical entities; in this view, you can understand that a pump doesn't suck anything, but rather that it isolates and compresses gas that moves into its input port, until (at a pressure greater than 1 atmosphere) it can be released to expand out of the pump's exhaust port.

Now you can understand the connections illustrated in Figure 11-1; note the special brass insert inside the ends of the two hoses, which are intended to couple to the fittings at the Reference-Pressure and Optical-Cell ports of the manifold.

You can intuit which valves you need to open to remove gas from your optical cell, and which valves need to stay closed to make this possible. To see if any of this is working, you'll want to hook up some pressure-measuring instrumentation to your system.

One instrument that your forepump system might already include is a thermocouple gauge. This will be of no high accuracy, precision, or linearity, but it has the advantage of measuring approximately the *absolute* (not the relative) pressure of the foreline, and of being sensitive in the 10 mTorr - 1 Torr range where the limiting capability of your forepump might lie.

The TeachSpin pressure transducer included with your interferometer has different capabilities: it is a symmetrical two-port device, and it measures only the *relative* pressure between its two ports. It delivers an electronic output of fine linearity, with a sensitivity very near ±10 V output for pressure differences of ±1.0 atmosphere between its ports.

It's easy to use one of the shorter hoses sent along with the transducer to connect one of its two ports, via a gas-handing T-connecter, directly to the cell. A longer hose can connect the T to the 'Opt Cell' port of the manifold; now one side of the pressure transducer is always subjected to the same pressure as exists inside the cell. To measure pressure changes in the cell, it is at first easiest to measure the cell's internal gas pressure relative to your lab's ambient air pressure, simply by leaving the second port of the pressure transducer open to air. [Later, you may choose to measure the gas cell's internal pressure relative to near-vacuum, by connecting the pressure transducer's second port to the 'Ref Pres' port on the manifold, and using the valves and forepump to establish and maintain a good low pressure there.]

Now as you remove gas from the cell, and (via the T) from one side of the transducer, you should see a transducer output changing from near-0 Volts toward +10 V (or -10 V, depending on which port of the transducer you've exposed to the cell). Only when you're 99.9% of the way to complete vacuum in your system will your thermocouple gauge even start to register, i.e. move off the top end of its scale, where all pressures from 1 - 760 Torr are crammed together into the same place. If you can achieve pressures under 100 mTorr in your cell, your vacuum system is certainly adequate for all that you'll be doing interferometrically.

In particular, if you know the absolute atmospheric pressure in your lab, as read by a barometer or manometer, and if you know the absolute pressure in your cell, as indicated by the thermocouple gauge in your system, then you know the pressure difference on your TeachSpin transducer in absolute units, and can calibrate its sensitivity $dV_{out}/d(\Delta p)$.

You still have a port, and a valve, on the manifold by which you can permit some ambient air, or some other supplied gas, to leak back into your otherwise isolated gas cell. Your transducer will show you how much pressure change you are causing, and your interferometer will show you the interferometric consequences. You should be able to control the inflow valve carefully enough to 'park' the cell pressure at any chosen pressure intermediate between vacuum and 1 atmosphere.

If you are using a compressed-gas tank for the supply of gas to fill the cell, be sure that your tank's regulator is set to deliver gas I above 1 atm of (absolute) pressure, i.e. that it's set to *well* under 1 atmosphere of gauge pressure. That's to ensure that you don't get an absolute pressure above (say) 1.1 atm inside your cell.

Once you've taken data along the route of filling the cell up from vacuum, you can take more data along the route of pumping the cell down to vacuum again; the trick is to open either the valve from gas supply to cell, or the valve from cell to vacuum pump, but not both at once.

When you want to change form one gas to another inside the cell, you need to pump away the remnants of your old gas, down to near vacuum, and then fill up with your new gas, *and* perhaps repeat this fill-and-empty cycle a few times. Why? To ensure, by successive, nearly complete, extractions of the gas in the cell, that all of the old gas is actually removed. If the 'old gas' was air, and the 'new gas' is helium, you'll be interferometrically quite sensitive to the effects of any residual air contaminating your otherwise pure helium.

There's finally a 'bleed' valve in place on your manifold, which you've kept closed for all the manipulations above. The purpose of this valve is to let air back into the whole system on the occasion that you turn off your forepump, and this 'back-filling operation' is motivated by the dire consequences that can occur if a forepump is turned off with a pressure difference left across it. ['Suck-back' is the baleful phrase, and an entire gas system filled with sticky pump oil is the possible consequence.] The right procedure is to open the bleed valve and the pump valve while the pump is still running, such that you hear air hissing into the bleed valve, and being pumped by the forepump. Now with air in flow, you can turn off the pump, leaving the bleed and pump valves open; that procedure ensures that the pump will settle down to a condition of zero pressure difference across it.