CHAPTER 12

Building Optical Systems

Here is Edward Bear, coming downstairs now, bump, bump, bump on the back of his head, behind Christopher Robin. It is, as far as he knows, the only way of coming downstairs, but sometimes he feels there really is another way, if only he could stop bumping for a moment and think of it.

—A. A. Milne, *Winnie The Pooh*

12.1 INTRODUCTION

Once you have a measurement principle, a photon budget, and a signal processing strategy for your measurement (at least in rough), you'll want to try it out, good and hard. Don't chicken out on this—the best failures are the ones that happen early. The first prototype serves to test your assumptions, not only about the physics involved but also the suitability of your strategy when faced with real-world components and situations. Don't tiptoe around the prototype; if the instrument is a big success, it'll have to work well without you around, and knowing what vulnerabilities you have to overcome in the real design is extremely valuable. Bang on it.

Plan your experimental approach to eliminate the most egregious technical risks first—the ones that can reduce your project to rubble. Things such as whether any affordable diode laser can reliably be made to simultaneously current-tune, frequency-modulate, and remain stable. Or whether you really can get the nonuniformity, readout noise, and dark current of your CCD detector down low enough without cryogenic cooling. Or whether a mass produced molded optical system can have its aberrations custom-adjusted in a stable way by indenting a TIR surface with a ball bearing on an XYZ stage.

After the first prototype has set your mind at rest, modify it or build another one to demonstrate end-to-end function and (a bit later) performance. Don't break it in the process—it is important to have a system functioning at some level all the time. This helps with sales pitches but, more importantly, keeps you from going too far astray if you make a design mistake. We've seen that an electro-optical system will go beyond our range of competence often, so this sort of belt-and-suspenders approach is vital to prevent disasters.

This chapter concentrates on how to go about assembling and aligning prototypes and limited-production optical systems in instruments. High volume optomechanical engineering is a specialized field, and if you're planning on building more than a dozen units, it is very worthwhile to consult someone who's good at it—it's not an amateur's job.

Building Electro-Optical Systems, Making It All Work, Second Edition, By Philip C. D. Hobbs

12.2 BUILD WHAT YOU DESIGNED

There is something about building a prototype system on an optical table or breadboard that sometimes seems to disable the higher functions of a designer's brain. After spending three days on the detailed Fourier optics of the system, he might accept a 10% vignetting of a Gaussian beam by a poorly placed prism; or fail to notice strong scatter brightening up the vicinity of his photodiodes; or carefully adjust and cement a diode laser collimator without making sure the beam axis is really horizontal.[†]

Since optical theory works pretty well, these examples of catastrophic synapse failure have the unpleasant consequences we'd expect. The first imperative in building an instrument is to really build what you designed. If the beam isn't clean, find out why. If it gets vignetted, move it. If there's a lot of stray light around, you've probably got a dirty element, a wayward beam hitting a ground glass surface, or a poor quality component with a lot of surface or internal scatter. It's in there, and it's doing you no good. Find it.

Debugging an optical system is an onion problem: you peel off a layer, cry, and peel off the next. A sufficiently hard debugging problem equals a failure. Being compulsive about tracking down little unexpected things you notice will be richly rewarded: your expertise with the system will grow more rapidly, and it will have fewer flaky problems. Debugging therefore gets easier quadratically.

12.3 ASSEMBLING LAB SYSTEMS

Most lab systems are built out of a mixture of purchased mounts, custom-machined pieces, and random chunks of metal held together with glue, double-sticky tape, and a few bolts. As a prototype gets modified more and more, the nice mounts deteriorate and there are more and more glued things, which makes the prototype get flakier and flakier. This is the normal life cycle. The key is not to break too many expensive optical parts when the glue fails under normal wear and tear.

12.3.1 Build Horizontally

Even if your instrument will be used with the optical path predominantly vertical (e.g., a microscope), build the prototype horizontally. It'll be much more stable and much easier to assemble and align. Also—crucially—it won't need so many custom-machined parts. All laser systems are interferometers, usually unintentionally. There will almost certainly be stray fringes there that will make life even more difficult if you build a floppy system.

12.3.2 Use Metal

Plastic is often a lot easier to work with than metal, being attractive, flexible, easily cut, and corrosion resistant. These attributes make it great for boxes, thin shims, baffles, and shrouds. Unfortunately, plastics have horrible dimensional instability. Your average plastic has 10× the CTE of ordinary metal, spreads heat 500 times more slowly (so that temperature changes lead to ∼ 5000× more warping). It also creeps under load and absorbs enough moisture to make it swell by a percent or so. Some plastics that are tops

[†]The author is not immune: that's where the examples came from.

in other ways (e.g., PTFE) have the worst thermal and mechanical properties. If you use plastics for optical mounts, be sure you understand, quantify, and accommodate these effects. Soft sheet aluminum (1100-T0) is almost as easy to work with and has none of these problems.

12.3.3 Scribble on the Optical Table

We often need to remove and replace some mount or stage, to an accuracy of better than 1 mm. This is difficult on optical tables and breadboards, which usually use 1/4-20 bolts in plain holes. To avoid problems, do one of two things: build a corner out of two post flanges for the mount or stage to snuggle up into, or mark the optical table by running a scriber or sharp-tip felt pen around the perimeter of the mount. You can easily do 0.05 mm this way, and especially with mag bases or universal mounts (the ones that can twist and translate), it saves a great deal of irritation.

12.3.4 Mounts

All the manufacturers' catalogs are full of beautiful optical mounts, translation stages, and so on; you can get six-axis adjustable stages with differential-screw micrometers and amazing specifications, if you're rich enough. You can also get snookered by specsmanship. A six-axis stage is about as stiff as overcooked broccoli, and a so-called kinematic two-tilt mirror mount exhibits hysteresis, wobble, nonorthogonality, and serious stick–slip energy storage[†] (see Section 12.9.2). Your best bet is to use fixed mounts for most things and keep the adjustable ones for where they're really needed. Whatever you do, don't believe translation stage manufacturers when they claim 0.1 μm stability and repeatability; they may have seen that in the factory, but for us it's more like 0.5 μm on a good day, with a new stage and a light load. Bang the stage just once, and the resolution becomes 2 μm, forever.

The basic trouble with most mounts is that they have a huge adjustment range and attempt to get precision and stability by using a fine thread. This is okay for some things, but if you think you can get a stable setup this way, think again.

You can get flexible shafts (made of wire rope) that connect the adjustment screws to handwheels bolted to the table, and these certainly reduce the amount of incidental motion due to hand pressure on the mount. They do nothing to reduce the mount's inherent floppiness, however; for that, you need a good design using verniers instead of coarse tweaks with fine screws.

12.3.5 Use Microbench for Complicated Systems

Individual elements in individual mounts can be placed anywhere you like. That's the good news and also the bad news. If you need several elements to be aligned correctly (e.g., a beam expander), use Microbench[‡] or the equivalent—systems based on mounts that thread like beads on centerless-ground steel rods. Their performance is quite good, and they allow easy adjustment of focusing. A hidden benefit is that it becomes natural to use fixed mounts for almost everything.

[†]The phenomenon is familiar from earthquakes.
[‡]Manufactured by Spindler & Hoyer.

There are a few problems with Microbench, too. The enforced alignment encourages etalon fringes and makes it difficult to use beam paths at odd angles. Components mounted on the sides of the rails are not as well located, because they rely on screws in plain holes (which are not stabilized by the rods).

12.3.6 Machine a Base Plate

Systems using noncollimated beams, narrow-angle components (e.g., Brewster windows or Pockels cells), or off-axis mirrors are hard to align from scratch with mounts bolted to an optical table. It gets a lot easier if you can make a base plate on a milling machine. Choose enough strategically placed adjustments and nail all the rest still. A base plate can easily have shallow slots and corners milled into it to make rough alignment automatic and fine alignment easy. This does require calculating the optical layout, but that is not normally too onerous. If the system is going to require a lot of fiddling with, you can use dowel pins instead of slots and corners, but that needs precise work—tolerances of $0.0001''$ ($2.5\ \mu$m) are typical.

12.3.7 Use Irises

Aligning beams gets a lot easier if you have a few iris diaphragms sprinkled around. An iris on each end of the beam path makes it simple to get the beam centered and allows you to tighten the irises down to get more precise location. This is especially useful with micro-optic components in Microbench setups, where just finding where you are can be a huge pain.

12.3.8 Getting the Right Height

A perennial hassle with optical mounts is that they're never the right height. If you're using sliding posts, that isn't the end of the world, but whenever we need to align one subassembly with another, height is always a problem. About 15% of the work of assembling the optical system is getting it the right height, and the easy methods (e.g., lab jacks and z translation stages) aren't very stable or rigid. One certainly needs a decent stock of shims, but they're never quite right, and the amount of cutting and trying is astronomical. What to do?

Use a pad of sticky notes. They come apart and go back together again, unlike shim stock, and they're a pretty repeatable 0.1 mm thick each; the thickness of the glue is negligible. There's no other shim material that makes it so easy to adjust beam positions in 100 μm increments, and you can use a paper punch to make the mounting holes. For thick shims (say, over 0.5 mm), you'll usually need something more stable, but a couple of metal shims and a few residual stickies work fine. (If the mild sleaziness of this trick bothers you, trim the paper with a razor blade after assembly so nobody sees what you did.)

Aluminum soda pop cans are a bit thinner, 70–100 μm thick depending on what height on the can you measure. They are made of very hard aluminum, so they work well in places where sticky notes are unsuitable. Their large thickness nonuniformity and their tendency to pucker makes pop can shims unsuitable for stacking. The wedge-shaped nonuniformity can be a help, since you can slide the shim a little for fine adjustments.

12.3.9 Light-Tightness

A dark-field system, especially a photon counting or high-accuracy photometric one, has got to be light-tight. This does not happen without thought. A butt joint is never light-tight—at least two 90° bends between close-fitting black surfaces is required, as when two tubes are butted together inside a sleeve, and more is better. A really light-tight system performs just the same in bright sunlight as in darkness.

12.3.10 Chop Up 35 mm SLR Cameras

A 35 mm SLR camera is a surprisingly convenient and economical optical toolkit for many things. Its two best features are a through-the-lens viewfinder and a film plane, where you can just butt your image receiver (e.g., a fiber plate, CCD, spatial filter, or quad cell) and be sure that it's focused where the viewfinder says it is. We do have to overcome our natural revulsion at hacking beautiful things, but they're so cheap nowadays that we can buy one to look at and one to sacrifice (the manufacturer is in it for the money anyway).

Modern cameras usually don't have mirror lockup levers, and many don't even have T or B shutter settings, which allow the shutter to be kept open continuously. A strategically placed piece of music wire will allow you to raise and lower the mirror manually, and a milling machine or a Dremel tool will remove the shutter in a couple of minutes.

12.3.11 Try to Use at Least One Screw per Component

It's usually the small mounts stuck on the corner of some translation stage with Krazy Glue that get knocked off when your hand slips.† Optical mounts are coated with Tufram or some other smooth and durable anodized finish, which glue does not stick to very well. Don't hesitate to drill holes in commercial mounts; they are comparatively cheap and will survive most kinds of abuse apart from banging them and brinelling the bearings. Use a small cordless hand drill and hand taps to put threaded holes in odd places on mounts, then use the glue to keep the mount from rotating. Setups built this way survive better, and the mounts and stages actually accumulate less damage, since you don't have to scrape the anodizing off to glue them.

12.3.12 Detector Alignment Needs Thought

Aligning dimly visible laser beams on small detectors is enough to drive you nuts. The surface is specular, so you can't go by the reflection unless you have a strong visible beam. The detectors are snuggled right up into the optical system, so you can't get a good view of them anyway. The best defense is to use the accuracy of the milling machine to do your work for you. Position the detector accurately in its module, position the beam accurately, then drop the module in (preferably located using roll pins or taper pins). You're done.

A good method for measuring the beam position is to mount an eyepiece reticle (available inexpensively, e.g., from Edmund) in the same type of mount you're using for the detector, and point a video camera plus macro lens at it. Put a thin coat of white

†Written an hour after destroying a $400 diffraction grating this way.

spray paint on the ruled side of the reticle to make the beams show up and eliminate parallax error. Shine the beams on the painted surface and look from the opposite side, where there is a nice bright dot with a scale superimposed. A frame grabber will make a permanent record for your lab book.

It's harder with multiple detectors, for example a differential system with two photodiodes and a Wollaston prism. One good solution is to put the detector die right in the plane of the top of the box, or embed it in a small grounded plate whose top surface is coplanar with the die. The benefit of this is that you can slide the detector module off to one side by some known amount, adjust the beam on a target on the box where it's nice and visible, and then slide the detector back again. You can draw the target under a binocular microscope with an eyepiece reticle. It doesn't get any easier than that.

A viewer that can be arranged to look at the photodiodes near normal incidence (and with adequate magnification) is even better, but normally much harder to arrange; if there's space, a clean microscope cover glass at an angle, with a camera looking at the reflection of the diodes, can be a godsend; the cover glass is too thin (100 μm) to offset the beams by much. Watch out for the fringes, though, and don't leave it in there when you're done. Another alternative is a mockup detector module with a piece of ground glass with two dots of opaque where the centers of the diodes lie. In either case, you'll need to put an external alignment reference on the module, e.g. a couple of roll pins or (for bigger detectors) countersunk mounting screws, which are mildly self-locating. It is really worth the investment of design and machining time to get the detector alignment right.

Aside: Photographic Paper. In the visible and UV, you can measure the positions of beams in inaccessible spots by using a chunk of photographic printing paper. You don't develop it or anything; just stick it in and turn on the beam for a moment. The resulting black spots are stable under room lights for at least several minutes, easily long enough for you to do your measurement with calipers or a measuring microscope. This is a great way to measure the separation of two spots for two detectors, or to measure the position of the beam in the clear aperture when you can't get at it to use the white-painted reticle trick.

12.3.13 Do-It-Yourself Spatial Filters

Making decent pinholes yourself is easy if you have an excimer or femtosecond laser handy to drill them. If you aren't so fortunate, you can use a trick: cut several small sheets of aluminum foil, make them into a stack, and stick a sewing needle partway through the stack—through the top layer, but not all the way out the other side. When you take the stack apart, you'll have a series of smaller and smaller holes; the smallest one will probably be useful. Measure the diameter by shining a (collimated) laser beam on it, finding the angular diameter of the first diffraction null and using (9.36). You can also do it with an evaporator and some polystyrene latex spheres: apply a (very dilute) aerosol of the spheres to a glass surface, and then put down 25 Å of chromium and ~250 Å of gold or aluminum by directional evaporation, and wash off the spheres. This gives you very accurate diameters and clean edges, but zero control over pinhole position.

You can easily make good slits in the lab from single-edged razor blades. Their edges have a 1 mm wide bevel, which is cut at a big enough angle to avoid reflection back into a laser, and they're very straight. Utility knife blades (the big trapezoidal ones) have

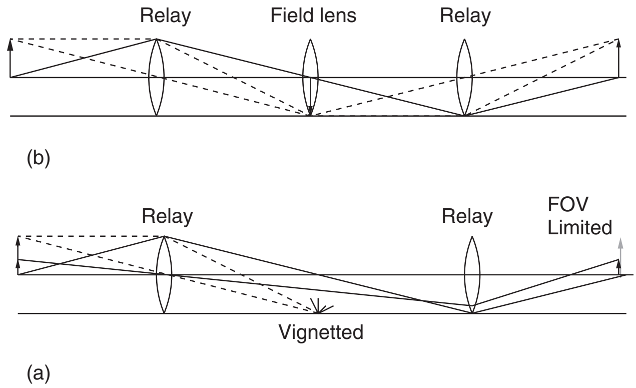

Figure 12.1. A field lens placed at an intermediate image doesn't change the image position but increases the field of view by corralling rays that would otherwise be vignetted. (a) Double relay lens has limited FOV. (b) Adding a field lens roughly doubles the FOV diameter without requiring a larger diameter system.

wider bevels for broader beams, and shaving razor blades are very smooth, although fiddly to use since you have to disassemble the plastic cartridge.

The easy way to make slit assemblies is to take a piece of plastic magnetic strip (as used in refrigerator magnets, with glue on one side), and drill a hole in it for the beam. Arrange the razor blades on the strip, adjust them until they're just right (perhaps with a shim), and then put on a drop of glue to secure them. This trick is great for tight spaces, and since you can stack several blades, it's easy to make weird shapes. Silicon makes atomically sharp cleaves in the [111] direction, so you can make good sharp edges by cleaving.

It's often convenient to use a slit with a slight V in it, made by crossing a pair of razor blades at a very small angle. This is useful with round beams, where the V allows adjustment of the slit width by moving it up and down, but doesn't affect the filtering operation much. You obviously wouldn't do this in a grating spectrometer, where the whole length of the slit is used at once.

12.3.14 Field Lenses

Optical detectors are not phase sensitive, so we can put an arbitrary phase screen at an image without damaging it. If we put a lens there, it doesn't change the image intensity distribution, but it does redirect the rays, as shown in Figure 12.1. This helps preserve the field of view of optical systems without requiring absurdly large diameter elements. The same idea, redirecting rays without altering the image, is used in Section 5.7.9 to improve rear projection screens.

12.4 ALIGNMENT AND TESTING

Testing optical components to verify their performance is a big field, and there isn't space to do it justice here. Instrument builders usually buy their lenses and mirrors, so exact measurement of figure errors is the vendor's problem. What we need to be able to do is to verify that a particular component is doing its job, align a complex optical system, and make sure our beam or image quality is what we expect. Malacara and colleagues give a very extensive discussion of many kinds of optical tests (which portions of this section follow), plus good references, and Yoder talks a lot about how to design mounts and align optics within a complex lens. The old-but-good MIL-HDBK-141 (Optical Design) covers much of the same ground and is a free download. (See the Appendix for these and other references.) We won't go beyond qualitative tests of the sort that are useful in debugging and troubleshooting, and instead concentrate on alignment.

12.4.1 Things to Count On

Considerations like the following underlie all the simple alignment schemes we'll consider here:

1. Light in a uniform medium travels in straight lines.
2. Lines parallel to a given line are parallel to each other.
3. Making a beam retrace its path exactly guarantees that it was normally incident (the *autocollimation* condition).
4. Rays passing through the center of a lens are undeviated in angle (although they are displaced slightly due to thickness of the lens).
5. A spherical surface rotated eccentrically will make a beam wobble around.
6. An accurate spherical surface is easy to make, and once it exists, it can be located very accurately by shoving it up against a well-turned circular mounting ring. As long as the optical surface stays in contact with the ring all round, the center of the sphere will lie on the axis of the circle.
7. A lens element has two surfaces. For any two nonconcentric spheres, there is exactly one line that joins their centers, and this is the lens axis. That means that within our tolerances for thickness variation in the finished lens, we can polish the two surfaces more or less independently and still wind up with the right lens when we're all done.
8. A beam that hits the vertex of an accurate corner cube is autocollimated.

12.4.2 Clamping

One consequence of facts 6 and 7 is that a lens barrel or mounting cell of high enough precision can automatically center a lens element, by forcing both surfaces into intimate contact with two circular rings lying in parallel planes, as shown in Figure 12.2. Since the two circles are coaxial, the centers of both spherical surfaces are forced to lie on the axis too. Yoder says that this produces adequate centration as long as

$$Z = \left| \frac{Y_1}{2R_1} - \frac{Y_2}{2R_2} \right| > 0.07, \qquad (12.1)$$

where Y_i is the radius of the ith clamp ring (remember that $R_2 < 0$ for a biconvex lens). Thus a lens element can usually be adequately centered by clamping it on its optical surfaces.

It is not altogether trivial to get the two rings aligned sufficiently well, especially if one is located only by its threaded edge; a high precision running fit of one ring into the

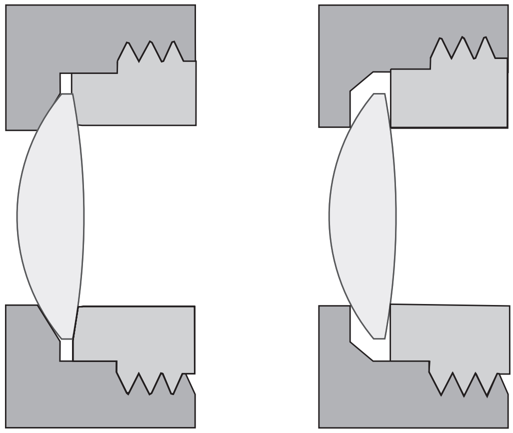

Figure 12.2. Clamping between two parallel circular rings enforces centration.

other works better. If you secure the ring with a threaded section, make sure the threads are loose, or they may bind due to overconstraint.

12.4.3 Soft Lenses

For soft or friable lens materials (e.g., ZnSe or NaCl), consider using an indium gasket. Indium is a very soft metal that has reasonable thermal conductivity and is available in wire form. Cut a length of indium wire slightly less than the circumference of the lens, place it in the fixed side of the mount, and put the lens in against it. A spring-loaded locating ring will be needed on the other side. Over several hours to a few days, the indium will deform to fit the surface of the lens, so that it will be held in without serious stress or deformation. This is especially helpful in high power CO_2 laser setups, where the high thermal conductivity of the indium and the excellent thermal joint between it and the lens improve the lens cooling dramatically. The bad news is that the centering action is much less since the lens has some freedom to rotate out of the plane of the mounting ring on the indium side.

12.4.4 Dimensional Stability

Clamping is a good example of a dimensionally stable mounting method; as long as thermal expansion doesn't cause the lens to lose contact with the ring, it'll stay still. When using glued joints, shrinkage causes the parts to move as the glue sets, and stability is limited by the thermal expansion of the glue. This can be dealt with by keeping the glue layer very thin, of course, but if there's a CTE mismatch, thin glue layers tend to delaminate due to extreme shear stress. Alternatively, you can load the glue with glass spheres of uniform diameter and apply a small preload while it cures. This forces the spheres into intimate contact with both surfaces (so the spacing will be set by the glass rather than the glue) while leaving the glue thick enough to handle the thermal strain. Choose glue with enough shrinkage that it will remain in tension over temperature—with 1% shrinkage and 10^{-4}/K CTE, it'll take 100°C change to make the glass spheres lose contact. With thermosets such as epoxy, a postcure bake will help as well—5 minute epoxy is horribly gooey and unstable, but an hour at 100°C improves it amazingly.

12.4.5 Beam Quality

In laser-based systems, we are constantly concerned with beam quality—but what do we mean by that, exactly? There are various measures in common use, such as beam divergence, M^2, Strehl ratio, RMS phase ripple, and so forth, and you'll see them all quoted. All these are static or time-averaged properties, which take no account of noise, let alone the spatial variation of the noise† (see Section 2.13). We're primarily interested in phase errors, which are what we mean by aberrations, but intensity nonuniformity is also important. The main thing is that if you want to use the fancy measurements, you need a fancy measuring interferometer that will give them to you, and if you don't have one, they're no use to you, except in deciding which lasers not to buy. Beam divergence is a partial exception, since you can measure that on the back wall of the lab for sufficiently small beams, but it isn't as useful since it depends sensitively on beam diameter and intensity profile.

As a practical matter, the two best lab tests of beam quality are a measuring interferometer for quantitative work,‡ and a combination of a dithered shear plate device (Section 12.6.4) for collimation plus a Foucault knife edge test for aberrations. For very small diameter beams (less than 3 mm in the visible), we can just shine the beam on the back wall of the lab and have a look.

Aside: Siegman's M^2. Many high powered lasers produce very ugly beams. Multistripe diode lasers, unstable resonators, and high pulse power beams don't have a nice single transverse mode structure, and because their different modes have different frequencies, they can have deceptively smooth near-field patterns that look a lot like Gaussians but have far worse divergence. In these cases, Siegman's M^2 beam propagation factor gives the best single-number summary of the beam propagation behavior. M^2 is the product of the second moments of the near-field pattern in (x, y) space and the far-field pattern in (u, v) space, normalized to 1 for a TEM_{00} Gaussian beam. Taking the second moment (i.e., multiplying by $x^2 + y^2$ and integrating over the xy plane) vastly overemphasizes small artifacts in the wings of the beam, so M^2 values near 1 aren't as useful as the Strehl ratio, which is why we haven't used M^2 much in this book. On the other hand, uglier beams can easily have M^2 of 10 or more, in which case M^2 is informative and Strehl is useless.

12.4.6 Image Quality

A laser beam has a single definite path, which normally doesn't change much except in preobjective scan systems. Imaging systems are expected to form images of any object presented to them, usually over a wide range of distance, field angle, and aperture, which makes image quality a much more elusive beast than good beam focusing.

The best measure of image quality depends on circumstances, but the most generally useful metric is the modulation transfer function as a function of position and orientation.

†By now, this should make you suspicious, but that's not the main point here.

‡The author fondly remembers an old Wyko Ladite that was hard to align but gave beautiful results. It got sent to the warehouse because no other machine could read its data.

As we saw in Section 9.3.13, the MTF is the modulus squared of the OTF, normalized to unity at 0 frequency (which automatically corrects for vignetting and power losses). The usual quick method for checking this uses a radial bar pattern with a circle drawn on it to show the nominal resolution limit. Moving this around the field gives a good idea of the lmiting resolution versus position. For a more accurate version, we use a set of sinusoidal amplitude gratings of different pitches and measure their contrast as a function of position, azimuth, and frequency.

12.5 OPTICAL ASSEMBLY AND ALIGNMENT PHILOSOPHY

There are two general classes of alignment test, based on beam steering and interference fringes. Beam steering tests are much easier but require lots of room.

12.5.1 Stability

The first requirement of any adjustment procedure is that it be *stable*. As the iteration proceeds, the corrections at each pass must be steadily decreasing. Stability can be achieved by performing measurements at each stage, or by choosing a design and a procedure such that the alignment is inherently stable. (Towing a trailer is inherently stable, but backing one up requires constant corrections to avoid failure.)

In cases where the adjustment is not iterative, for example, progressively adding components to an optical bench setup, we require in addition that our alignment criterion remain valid at all times during the assembly; it is not much help to put a HeNe alignment beam down the optical axis if the first lens is made of germanium, or to allow errors in placing the earlier elements to steer the beam completely out of the system by the time the later elements are added; the assembly procedure as well as the alignment procedure must be stable.

Since most optical adjustments are nonlinear in character, the easiest way to ensure stability is to make sure that the interaction between controls is small, and then proceed in the right order.

12.5.2 Orthogonality

Two adjustments are said to be *orthogonal* if changing one leaves the other unaffected; for example, x and y on a good translation stage, or translations in the object and translations in the pupil (i.e., position and angle). Orthogonality is the Holy Grail of the manufacturing department: set each screw once and forget it. An example is the beam steering system of Figure 11.2. It isn't always easy to achieve, but really work at it before giving up; a stable alignment strategy will prevent errors from growing, yet if the controls interact a lot your alignment algorithm may never converge. Consider aiming a beam at a spot on the wall using nonorthogonal mirror adjustments: if moving the spot 10 cm to the right needs 1-turn adjustment of one screw, then any accompanying up and down motion should require no more than 1/10 turn of the other screw to correct it. It gets a lot worse when there are fewer readouts than adjustments; and it's worst when there's only one, as in maximizing the output power of a laser.

A general rule for such nonlinear 1D adjustments is that an interaction of 1/10 turn between two knobs is pretty manageable, but 3/10 requires an expert and 5/10 is impossible to do by hand unless the system is nearly aligned to start with. This rule applies in both directions, so that it does not help to put one adjustment on an extremely fine thread. It's a bit easier in two or more dimensions, but the point remains: avoid interacting controls if at all possible.

12.5.3 Use Serendipitous Information

As we get more experienced with a system, we develop powerful intuition about it—we get faster at adjusting it, and we know what's wrong when it breaks. One of the main sources of this intuition is noticing the information that can be gained from unintended places. For instance, focusing a laser microscope gets a lot easier when there's an incidental beam somewhere whose diameter changes with focus; put a card there, and trace the beam outline on it when the focus is perfect. Forever afterwards, rough focusing will be a piece of cake. Diffraction patterns in stray return beams are similarly useful for finding features on a substrate. The colors of the white-light fringes in an interferometer are another good pointer. Be on the lookout for these in your systems; they'll save you lots of time and frustration.

12.6 COLLIMATING BEAMS

Collimated beams are very easy to work with, which makes them popular. Collimation is particularly easy to test for, so you can be sure when you're there. You can change the spacing between elements freely (at least those with no optical power), put in polarizing cubes, and so on, without causing any change in the system aberrations—at most, a slight change in magnification. They have several nasty features, however, such as maximal sensitivity to etalon fringes and increased danger of eye damage if the beam is powerful.

How to collimate a beam depends on what it looks like to begin with. A beam from an extended source such as a light bulb cannot be collimated well, because its spatial coherence is not good enough. Here we specialize to laser beams, and so assume that the beam is monochromatic and fully coherent.

A collimated laser beam is one whose phase fronts are perfectly planar, so any transverse slice of the beam will have a constant envelope phase. A simple calculus of variations argument shows that this minimizes the equivalent width of the far-field pattern for a given intensity distribution. Thus collimation can be measured in three general ways: directly, by looking at the behavior of the beam as it propagates; in the far field, by measuring its optical Fourier transform; or in the wave picture, by measuring the phase fronts with an interferometer. Which of these is best depends mainly on the size of the beam.

12.6.1 Direct Collimation

A beam N wavelengths across stays collimated for N beam diameters. Direct collimation is done in the obvious way: stick a mirror in the beam to shine it on the back wall of the lab. Walk back and forth, using an iris diaphragm (best) or a sheet of paper with a circle drawn on it to gauge when the distant beam is the same diameter as that exiting the collimating lens. With a wall 3 meters away, this procedure will provide diffraction-limited results for visible beams of <2 mm diameter, but gets rapidly less accurate as the beam gets broader.

You can mount a corner cube or a mirror on the wall to get twice the path length with much less travel. Note that this is susceptible to strong $1/f$ noise from people occasionally walking through the beam. When you do this, make sure you avoid sending the beam where it may encounter someone's head. Optical tables are usually at just the right height for escaping beams to hit the eyes of someone sitting down—like you, for instance. If you need to do it a lot, a piece of 3 inch plastic drain pipe with a corner cube at the other end can stretch all along one wall of the lab with perfect safety. (Leave an opening at the far end so you can align it without peering down the tube.)

12.6.2 Fizeau Wedges

A Fizeau wedge is a very narrow triangular air gap between two optical flats with a thin spacer on one side only. Under narrowband illumination, it produces a system of parallel fringes that are easily visible. Fizeau wedges are good for gross collimation, because a collimated beam produces straight fringes of equal spacing, but for fine work, or in the presence of aberrations, Fizeau wedges are nearly useless. The strong asymmetry of the pattern, and the necessity of spotting, say, a 1/10 cycle irregularity in spacing or warping of a fringe, make it impossible to do by eye. Using a CCD camera to look at the fringes would have a much better chance, since the lithographically defined sampling grid of the CCD will show up deviations from the predicted peak positions very clearly. (Cast the fringes directly on the CCD, or at least calibrate the geometric distortion of the lens if you do this.)

12.6.3 Shear Plates

A shear plate is a sort of Fourier transform of a Fizeau wedge: instead of interfering two waves offset in angle, it uses two offset in space, for example, by reflection from the front and rear surfaces of a tilted flat plate; the phase of the fringes is thus a finite difference of the phase front of the incoming beam. A collimated beam produces a uniform intensity that goes up and down as the path difference varies, slight defocus changes that to a system of parallel lines, and more complicated errors to a more complicated pattern. Shear plates are insensitive to phase curvature orthogonal to the shear direction, so you need two plates for a general beam.

12.6.4 Collimeter

The late lamented Blue Sky Collimeter was a scanned shear plate interferometer. They don't make it any more, but the function isn't that hard to reproduce. The relative phase of the sheared beams is scanned by turning the plate slightly, using a flexure powered by a bit of Nitinol wire (a shape-memory alloy, sold as "Muscle Wire" by robot hobby suppliers). This makes the fringes move back and forth at about 0.2 Hz, which makes them easy to see, and easy to tell apart from the envelope of the beam intensity. When the beam blinks on and off rather than exhibiting spatial fringes, it is collimated to high accuracy ($\lambda/50$ or even better). These devices are most useful with good quality beams, where it is obvious what to adjust to fix the problem. Heavily aberrated beams (e.g., astigmatic diode lasers) produce complicated fringe patterns with odd motions, which makes adjustment a more subtle business. A slower, wider range motion would be better, and if you're doing many, you really need a proper measuring interferometer.

12.7 FOCUSING

12.7.1 Autocollimation

The autocollimation condition is useful for focusing as well as for alignment. By symmetry, if there's a surface at the focus of a beam, aligned normal to the beam axis, the beam will retrace its path. Not only will the two beams be coaxial, but if the incident beam is collimated somewhere, the reflected one will be too. This reduces the focusing problem to collimation of the reflected beam, but good accuracy requires the incoming beam to be collimated very well, since any collimation error turns into a focus error. This method is good for quick coarse focusing if you know there's no vignetting, since the two beam diameters will be identical in focus.

12.7.2 Direct Viewing

Even if you don't use the 35 mm camera trick, a viewer can make focusing a snap. We talked about this in Section 11.8, but it's worth reemphasizing. Use a microscope trinocular head, a couple of eyepieces, and a beamsplitter to form an image of your focal plane, put a target there, and adjust the focus of the beam until the spot on the target reaches its minimum diameter. If the viewer has a higher NA than the beam, you can spot beam defects this way too. You can do this in reflection, right through the objective lens and back again, so that no additional high-NA optics are required. Be sure the beam power at the eyepieces is below 10 μW.

12.7.3 Foucault Knife Edge

The Foucault knife edge test is a good way of locating an aerial focus. As Figure 12.3 shows, it's very simple; put a razor blade on a translation stage and move it across the beam while looking at the transmitted light. If the razor is beyond the focus, the shadow moves the same way as the blade, and if it's inside the focus, it moves the other way (this is obvious from geometric optics). What may be a bit less obvious is that when the blade is right at focus, the transmitted beam just fades out and blinks off, with no obvious shadow motion. This is not so for an aberrated beam, and in fact the knife edge test is a sensitive qualitative test for aberrations. Near focus, an astigmatic beam either gets chopped off from the inside out, or the cut edge goes left to right, swings round to up and down, and then goes to right to left, depending on how astigmatic it is. If the knife edge pattern is funky-looking, you've got at least a wave of aberration. For smallish aberrations, the pattern gets steadily better looking as the aberration decreases, but if you have more than a couple of waves of astigmatism or coma, it takes a trained eye to know what to adjust.

12.7.4 Intensity-Based Chopping Tests

There are beam scanning instruments available that use a slit in a rotating drum to produce a 1D beam profile. Alternatively, a rotating chopper will give you the integral of the 1D beam profile (don't use a vibrating chopper, because the blade velocity is poorly controlled). These are great for finding the beam waist, but not much use for diagnostics; unless the beam is invisible, or the NA is high, a viewer is better and faster. Their major

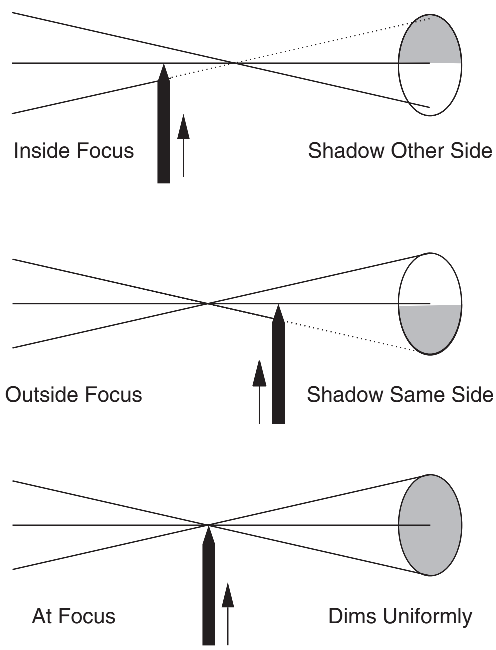

Figure 12.3. Knife edge test.

advantage over a direct viewer or a CCD camera is that they produce a trace for an oscilloscope, and they can work in the IR and UV. The other thing is that a chopper can be used as a reasonably accurate focus finder for good quality beams; with a split detector or quad cell in the beam, the output will be nearly 0 V all the time if the chopper is at focus, a positive pulse waveform for one side of focus, and a negative pulse for the other. It isn't perfect because the chopper blade is comparatively thick, so that the beam can bounce off the side as well as the top.

12.7.5 Diffraction Focusing

The focused image is the Fourier transform of the pupil distribution, so one good way of making sure your imaging detector (e.g., CCD) is at the focus of your optical system is to put a strong diffractor at the pupil and adjust the detector position until the Fraunhofer pattern is in sharp focus. The usual method is to put a wire or a piece of tape across your pupil, so that the Fraunhofer pattern exhibits a sharp line running normal to the wire. Adjust the CCD position until the pattern is a single sharp line (not two closely spaced lines), and you're in focus. This works best with long-focus systems focused on infinity, where the pupil position is not especially critical (it was first developed for astronomical telescopes).

12.7.6 Speckle Focusing

Unless it's very tightly focused, a laser beam bouncing off a rough sample (e.g., in a microscope) produces speckles in the scattered light. These are usually thought of as a noisy nuisance but (like other kinds of noise) can often be turned to advantage, as in speckle interferometry. One thing they're good for is focusing. If you shine the speckles on a card while moving the sample from side to side, they move opposite ways for samples on opposite sides of focus, and when the sample is in focus, they just blink on and off like sunlight through a tree's canopy.

12.7.7 Focusing Imagers

One good and very fast way to focus an image sensor is to use a point source at the correct distance, bright enough to cause blooming. Adjust the focus to maximize the length of the blooming tail in the image. This works because blooming is caused by charge overflowing from the CCD well, and the more light is concentrated on a small area, the more charge will overflow to cause blooming. This of course assumes that the CCD fill factor is high, or that the image of the point source is at least a few pixels in size, because otherwise if you have an unresolved point source that happens to be centered on an insensitive area, the blooming will *decrease* near focus instead.

If you have a repeatable, calibrated focus actuator, another approach is to take images at known focus positions, compute the contrast, run a low-order polynomial through the data points, and interpolate to get the position of best focus. The contrast can be computed as the sum over the image of the RMS values of the first finite differences in x and y. There are lots of variations on this; for an astronomical imager, you can look at the diameter of star images as a function of focus.

12.7.8 Standards

Film camera manufacturers have standardized the distance from the mounting flange to the film plane; each manufacturer has its own standard, unfortunately, but at least this has to be constant for all lenses in their lines. (Why?) If you have the matching camera, the film rails give you a fixed reference plane, and you can measure the focal position with calipers. That way, you can use dead reckoning for focus, at least with a sufficiently slow lens.

12.7.9 Subapertures

You can put a mask on top of your lens, consisting of two holes near opposite edges of the pupil, and adjust the focus until the focused spots coincide. This is a bit like a rangefinder camera and works very well. The hole diameters should be about one-quarter of the clear aperture so that the spots aren't too big.

12.8 ALIGNING BEAMS WITH OTHER BEAMS

It is frequently necessary to align two beams accurately with one another, propagating either in the same direction or opposite directions. For the copropagating case, we may want to wavelength-division multiplex a fiber link, or to pass two beams of slightly different frequency down exactly the same optical path so that any vibration or scan motion causes only common-mode errors. Counterpropagating examples include lining up your beam with a guide beam provided as an alignment aid, as in the Burleigh WaveMeter, or two-photon Doppler-free spectroscopy, where the resolution improvement requires the beams to be well aligned. How difficult this is really depends on whether it is enough for the beams to be aligned with each other, or whether they must in addition be aligned accurately with the rest of the optical system.

Check the overlap with a translucent screen, such as an index card or sticky note; by putting it in the beams, and then blocking and unblocking the brighter of the two, you can see the overlap very conveniently in both the co- and counterpropagating cases.

12.8.1 Copropagating Beams

Getting two beams to propagate together requires some sort of beam combiner, which is a beamsplitter run backwards, as shown in Figure 12.4a. Put the combiner on a two-axis tilt stage for one adjustment, and use a mirror on a two-axis tilt mount for the other. For applications where back-reflections are a serious concern, use a Wollaston prism for

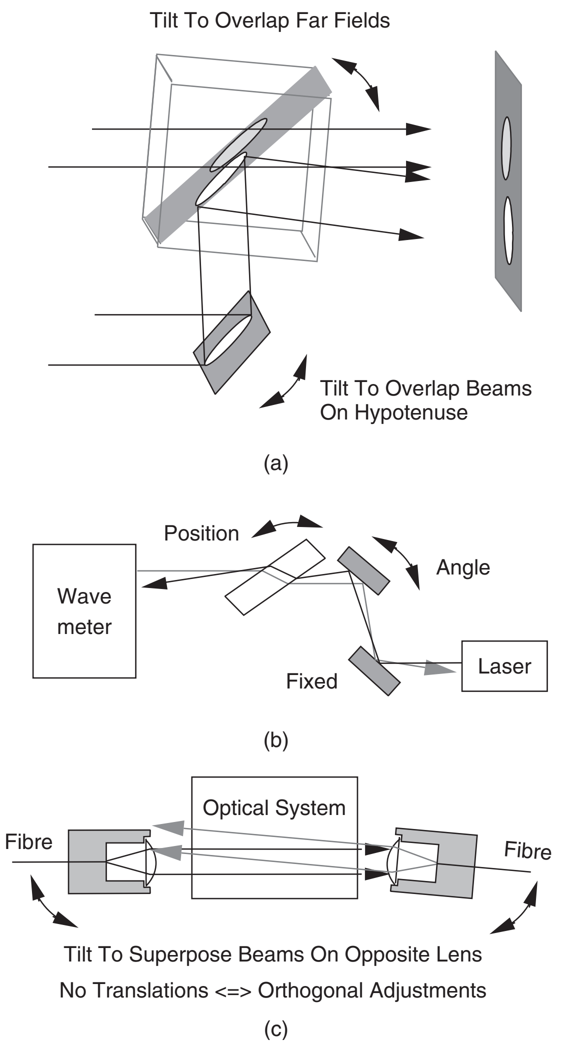

Figure 12.4. Aligning beams to other beams: three techniques.

the beam combiner.[†] Use the mirror to superimpose the beams on the combining surface and then tilt the combiner to align them in angle. Commercial mirror mounts and prism tables are relatively poor for this, because the adjustments are too tweaky, so consider using a stiff flexure for the combiner instead.

This is of course not a complete solution to the problem, because we often want to align two beams of very different wavelength, for example, a HeNe guide beam and a CO_2 cutting laser. You may have to use something a bit more special for the beam combiner, such as a hot or cold mirror, which passes one wavelength while reflecting the other. Watch out for subtle chromatic effects downstream of your combiner, especially polarization funnies.

Narrow beams can be aligned well by eye. A rule of thumb is that two beams of diameter d can be made coaxial to within the diffraction limit by eye on the first try, as long as the distance L between the two mirrors obeys

$$L > \frac{d^2}{5\lambda}. \tag{12.2}$$

This estimate takes account of the visual difficulty of aligning the centers of two smooth luminous spots, so variations among people limit its accuracy to $\pm 20\%$. You can make do with less space if the beams have strong features (e.g., diffraction rings or sharp edges) that can be used as visual cues. The main things to avoid are putting one beam steering element on a four-axis stage, and using two adjustments that are nearly degenerate, e.g. two tilting mirrors spaced very close together.

If the beams are too large for this, you can use a lens to Fourier transform the beams. Examine the focused spots with a magnifier or a low power microscope to check that the focused spots are really superimposed. Because the focused spot gets smaller as the beam gets broader, this technique takes care of the broad-beam case very well.

12.8.2 Constrained Beam-to-Beam Alignment

A slightly more difficult case is where the rest of the optical system provides more than one axis, which the beams must follow (Figure 12.4b shows the counterpropagating case), where, e.g. a laser and wavemeter each defines an axis, and we don't want to mess with the system alignment just to use the wavemeter. The figure shows one angular and one xy adjustment, which is not orthogonal but does converge rapidly because the xy alignment doesn't mess up the angular alignment.

If the beam alignment is critical, combine the beams as early as possible and pipe the results down the system so that subsequent pointing drifts affect both beams identically. An example is a common-path interferometer, where errors in the relative position of the beams cause an enormous degradation of the vibration immunity of the system. An orthogonal approach is to use the beam alignment jig of Figure 11.2, translating the beam and the jig together to make the beams overlap at the exit pupil of the jig, and then work the angular adjustment until they are perfectly aligned.

[†]It isn't so easy to use a Wollaston, because you can't just tilt it to change the alignment. You'll need another mirror.

12.8.3 Counterpropagating Beams

Aligning two counterpropagating beams is the easiest of all if you have lots of space, and the hardest otherwise. The best way to do it is to make both of them adjustable in angle, e.g. by bouncing them off mirrors on two axis mounts. Adjust Mirror 1 to overlap the beams on Mirror 2, then use Mirror 2 to overlap them on Mirror 1. (Figure 12.4c shows the same idea with fiber collimators.) These two adjustments are orthogonal, so the two beams should now be in alignment. If it is necessary to iterate, the mounts can be adjusted independently without interacting, which makes the alignment fast and convenient. Sometimes one beam cannot be moved, e.g. the one defining the optical axis of the system. Then the second best thing is to use one translation stage and one tilt stage. Adjusting the translator to overlap the beams on the tilting mirror and then adjusting the tilt to align the beams in angle is a workable two-step procedure with no interactions. If you need to iterate, it is no longer orthogonal. Every time you touch the translation stage, you have to readjust the tilt. The iterations are always stable, because the tilt does not misalign the translation, so that the iterations must terminate.

Both of these schemes work best when there's lots of room for angular misalignments to turn into spatial errors. If you have a single-mode fiber collimator with a beam coming out, a 3 meter spacing with a 3 mm HeNe beam will allow you to send another beam back into that fiber in one try. It gets harder from there down.

If there isn't enough space for the easy way, you have to use the harder ones, which rely on making the beams copropagating and using the tests of the previous section. You can do this in two ways: first, if it's okay to block the beam during alignment, you can use a beamsplitter and an accurate corner cube retroreflector, as in Figure 12.5. Second, if you need the beams unblocked during the test, you can put in a weak reflector (e.g., a microscope slide) and adjust it for autocollimation of the first beam. Leaving the first one fixed, align the second using the constrained copropagating technique. The first technique gives you the autocollimation free, which is considerably more convenient.

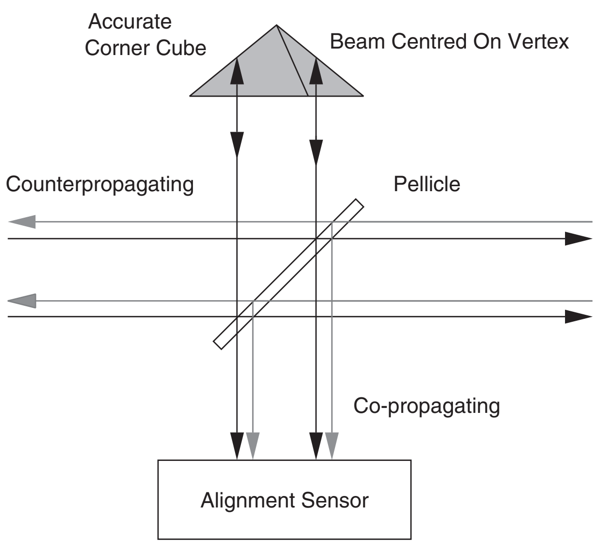

Figure 12.5. Alignment aid: corner cube and beamsplitter.

12.9 ADVANCED TWEAKING

12.9.1 Interferometers and Back-Reflections

Lasers are amazingly sensitive to back-reflections into their cavities. Michelson and Fabry–Perot interferometers are very bad for this, because they rely on back-reflection in order to operate. Mach–Zehnders are much better in this regard. Before aligning your system so that it just exactly blows your laser out of the water, it's important to think about what the results of perfect alignment will be. A lot of problems are reduced by strategic minor misalignment.

12.9.2 Backlash and Stick–Slip

Anyone who has ever successfully aligned a complicated laser system has educated fingers. There may be 20 knobs to twist, and the correct alignment setting is in there somewhere, but where? The problem is made worse by the usual tendency of the adjustments to exhibit twin evils: *backlash* and *stick–slip*. Backlash is like a loose tap handle in the shower: the same water temperature corresponds to different knob positions, depending on whether you're going clockwise or counterclockwise. Stick–slip is what keeps all the bolts in your car from falling out, but it's less useful in adjustment screws: when you twist the knob, it doesn't move smoothly, but goes in jerks, usually shaking all the other screws out of adjustment as it goes. Surfaces stick and slip because the coefficient of static friction is higher than the coefficient of dynamic friction. Energy stored in elastic deformation can be released as motion, causing instability over time—the San Andreas fault is a well-known example, but it happens in translation stages too.

The net effect is that even when you've located the region where the sweet spot lies, getting the alignment just right is an iterative procedure requiring patience.

Example 12.1: Aligning a Regenerative Amplifier. At this writing, the author's fanciest toy is a tunable picosecond laser source, producing 20 ps pulses continuously tunable from 420 nm to 10 μm, with a small hole around 710 nm. It's pumped by a frequency-tripled, flashlamp-pumped Nd:YAG laser that has both active (Pockels cell) and passive (dye) mode locking. It's a thing of great beauty but requires a lot of laser jocking to keep it working. Just the pump laser has three separate cavities. The most frequent adjustment is the alignment of the regenerative preamplifier section, which takes 5 μJ pulses and amplifies them to the 500 μJ that the power amplifier section expects.

The mirror adjustment has two stages per axis: ordinary ball screws for coarse setting, and a vernier consisting of a fine screw acting on an inclined plane, forcing the anvil opposite the screws in and out slightly. This vernier is a poor design that has a lot of stick–slip energy storage, requiring a lot of babying. To align a system like that, you have to do all your tweaking in the screw-tightening direction. The reason is that the screw is so much stiffer than the return spring that the same energy release corresponds to a far smaller motion, and hence much less instability. During tweaking, this also makes the X axis adjustment stay still while you're tweaking Y, which is a huge help.

Tweaking in one direction requires a bit of memory and judgment: for each axis in turn, slack the screw off until you're well out of alignment in that direction. Slowly tighten it until it starts to lase, and watch the energy meter to see what the highest value is. You have to tune past it to find this out, so on the second pass, tweak it near the peak position and then very slowly tighten until it reaches peak output, and stop. You need to remove your hand from the knob after each small motion, so this may take a few tries. On this particular laser system, the backlash is wider than the whole range in which lasing occurs, so it really matters.

12.9.3 Adding Verniers

If you have a system with a vernier (coarse–fine) adjustment, you may find that there are places where the coarse adjustment isn't stable enough for the vernier to be used. In that case, move the coarse adjustment somewhere nearby and try again. Verniers are far superior to just putting a fine screw on a coarse adjustment—if you're designing the system, make the vernier separate, for example, a very weak lens that you can shift around.

12.9.4 Cavities with Obstructions

When you're aligning an optical cavity containing obstructions such as pinholes, it is sometimes better to align the mirrors without the pinhole, then replace the pinhole. An obstruction produces such strong slopes on the merit function (e.g., power output) that the weaker ones due to the cavity alignment are often obscured. If you can't take out the pinhole, make sure you measure the sensitivity of the merit function versus pinhole position: align the mirrors, measure, misalign the pinhole slightly upwards, realign, measure, move the pinhole slightly down from its original position, realign, measure. Then you'll know whether to move the pinhole up or down to approach the optimum. Repeat in the horizontal plane and iterate to convergence. This is much more time-consuming than just taking out the pinhole. Mark the table so it's easy to put the pinhole back where you got it.

12.9.5 Aligning Two-Beam Interferometers

There are two kinds of interferometers: two-beam ones, such as Michelson and Mach–Zehnder types (see Section 1.6), and multibeam ones such as Fabry–Perot and Fizeau types. It is seldom enough to get the beams propagating in the same direction. They must also be in the same state of collimation and (often) have zero path difference between them. To do this, use the fringes between the two beams as a long-path shear plate.

Zero path difference is easily found using the white-light fringes. Thermal sources have coherence functions with a single peak (see Section 2.5.4), so fringes appear only near zero path difference. Make sure you cancel out the dispersion as in Figure 1.10: use a compensation plate in a Michelson and put the beamsplitters the right way round in a Mach–Zehnder. Otherwise, the different paths in glass produce dispersion that will wash out the white-light fringes and generally make life difficult. You can tell when you're approaching the white-light fringes by looking at the pattern through an interference filter—narrowing the bandwidth of the detection has the same effect on the fringe contrast as narrowing the bandwidth of the source. Thus a 50 nm bandwidth filter will allow you to see fringes several times further away from ZPD, which makes alignment easier. Use a single-point detector and an oscilloscope to measure fringe contrast for interferometer fine alignment, especially in the IR, where it's inconvenient to use the fringe pattern directly. It gives you a single parameter to optimize, which is easier to do when you get near the end. Mild mechanical vibration (e.g., gently pounding the table with your fist) will often give you moving fringes, which can be detected and displayed on an oscilloscope. Make sure you measure the photocurrent of both beams separately, and calculate the maximum contrast in advance using (12.3)—otherwise you don't know when you're there. The maximum fringe contrast available is

$$C_{\max} = \frac{2\sqrt{I_{\text{sig}} I_{\text{LO}}}}{I_{\text{sig}} + I_{\text{LO}}}, \tag{12.3}$$

which is 100% when $I_{\text{sig}} = I_{\text{LO}}$, but only 2% when $I_{\text{sig}} = 10^{-4} I_{\text{LO}}$.

A contrast of 100% makes the beams blink on and off, and it's a pretty sight to see an interferometer that good. Nonetheless, do remember that beams blinking on and off don't necessarily mean that your beams are in the *right* state of collimation, just the *same* state (much time has been wasted by confusing the two).

12.9.6 Measuring Focal Lengths

This effect can be very useful in precision focal length measurements—throw together a Michelson interferometer using an expanded laser beam, put a lens in one arm, and move the lens axially until the fringes blink on and off with no rings. The lens is now exactly one focal length from the mirror. (You'll want to put the lens in its cell first—otherwise the precision will be lost as soon as you take it out of the setup.)

12.9.7 Aligning Fabry–Perot Interferometers

There are two main ways of aligning plane mirror Fabry–Perots: either by using a laser (preferably one whose wavelength is on the skirts of the HR mirror coating) and make all the beams overlap, or put a large area line source behind it (e.g., a green mercury lamp) and make the ring spacing stay the same as you move your head from side to side. (You have to be looking from the right distance so that the fringes have a reasonable width.) Both of these are pretty sensitive and will get you into the ballpark very rapidly. Fine alignment can be done by maximizing the transmitted intensity of an expanded HeNe beam or some other appropriately isolated laser source.

12.9.8 Aligning Lasers

A laser is a special case of a Fabry–Perot, since no light comes out until it's aligned. After it starts to lase, fine adjustment is a matter of using an analog power meter and fiddling until it reaches a maximum; it's the initial coarse alignment that's the trick.

Since a laser is a Fabry–Perot after all, you can align it the same way, by using an auxiliary laser and aligning the resulting beams. Sometimes you have to do that, for example, when the laser is first manufactured and the resonator isn't focused, but usually it can be avoided in the lab.

There's another technique that is quicker and much more fun: raster-scanning the alignment with a screwdriver. At least one mirror will be sitting on a spring-loaded mount with adjustment screws. Slack off one screw all the way, and jam a pry bar (e.g., a big screwdriver) in next to it. Crank that one quickly back and forth through its range while slowly moving the other screw, starting from one end of its range. You'll get bright flashes when the slow adjustment is roughly right, so then you pull out the screwdriver and twist the fast knob until you get steady output. This quick method is great for *walking in* a gas laser after changing lines or cleaning the Brewster windows. Even though it's what laser service technicians do, it's a bit shocking to those who haven't seen it done before—almost as if you were using a ball-peen hammer and a tire iron—which makes it even more fun.

12.9.9 Aligning Spatial Filters

One of the reasons spatial filters aren't used more often is that they are fiddly to align if you don't do it properly. There are two phases: finding the beam and finding the sweet spot. Finding the beam is done the same way you align a laser but is a bit easier since the bright light is there all the time, rather than only when you're nearly aligned. If your pinhole is in a shiny surface, the easiest way to do rough alignment is to put in a viewer, even just pellicle or a microscope cover glass at 45° in the incoming beam, shining the reflected light onto a card. Defocus far enough that you can see the image of the pinhole on the card, then zero in by translating the pinhole until it's centered, adjusting focus, centering again, and so on until the pinhole nearly fills the image. By this point, light should be coming through your pinhole. If there is too much interaction between focus and lateral adjustments, your beam is misaligned. If this is not on purpose (remember that slight misalignment is a good idea when using a laser), adjust the lens laterally while racking focus in and out to eliminate the lateral motion of the pinhole on the card.

Fine alignment depends on whether your pinhole is resolved or not. If it isn't, just adjust the pinhole while looking at the detected photocurrent on a meter (analog ones make this a great deal faster). Find the maximum and you're done.

Wider pinholes or slits are a bit subtler. If you don't know which way to go to find focus, look at the far-field pattern of the beam (transmitted or reflected). Passing through focus reverses the direction in which the image of the pinhole moves when the object moves (this is obvious from geometric optics, since all the rays pass through the focus). In the transmitted beam, a pinhole too close to the lens makes an image that moves in the opposite direction. The directions are reversed in the reflected image.

12.9.10 Use Corner Cubes and Pentaprisms

We typically do alignment by overlapping spots in the far field and near field or, alternatively, by looking for interference fringes. Overlapping beams is a lot easier, especially for setup, but requires two measurements per iteration and needs lots of space or auxiliary aids. The most useful alignment aids are those based on accurate corner cubes and pentaprisms, which allow really good angular alignment of counterpropagating beams. You'll need an independent method for aligning them in position. If your cube has sharp enough corners, and your beam is big enough, you can bounce it right off the vertex (grey lines), in which case all is easy; a quad cell will tell you when they're coaxial.

If you can't do that, you'll have to let the returning beam be offset and align the two beams accurately parallel, using a ruler or (for very fine alignment) two quad cells spaced appropriately. You can do the same sort of thing with other constant-deviation prisms (e.g., pentaprisms).

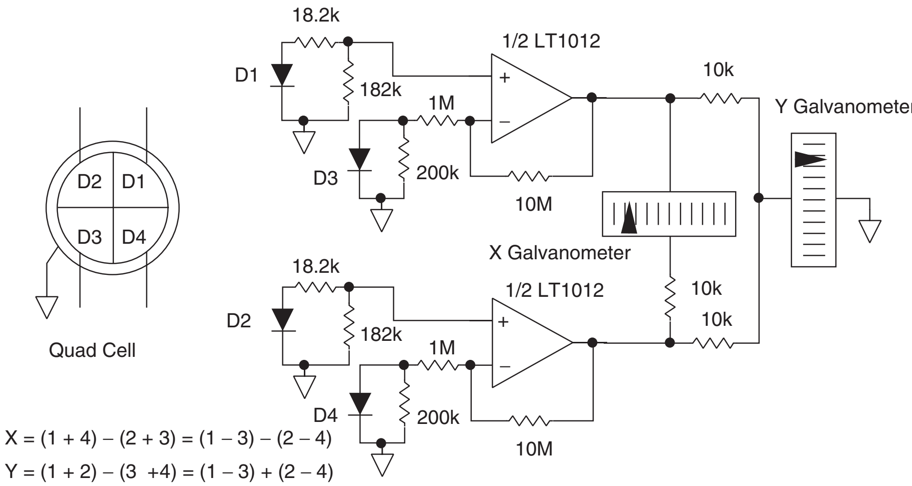

Figure 12.6. Quad cell alignment aid.

12.9.11 Use Quad Cells for XY Alignment

Another approach to aligning copropagating beams is to use a quadrant photodiode and a pair of analog voltmeters, as shown in Figure 12.6. Put a large (22 MΩ) resistor from each segment to the common terminal, and look at the combinations $(1 + 2)–(3 + 4)$ for the X direction and $(1 + 4)–(2 + 3)$ for Y. You can do all this with a dual op amp and a few resistors, as shown in Figure 12.6 (the trick is to notice that the X and Y signals are given by $(1 - 3) \pm (2 - 4)$). Loading of the open circuit photodiodes by the resistors will start to dominate at low photocurrents, reducing the signal range at the meters. A fancier version, using two high input impedance differential amps (maybe implemented with a quad op amp IC) instead of resistors and a dual op amp, will extend the logarithmic range to photocurrents of 5 μA and below without hurting the strong signal performance.

The reason this works so well is that the open circuit voltage V_{OC} of a photodiode is proportional to the logarithm of the optical power incident on it (see Section 14.6.1). For matched diodes, subtracting the logarithmic signals gives the logarithm of the ratio of the powers on each diode pair, and that ratio is independent of overall power. Thus we have a very simple meter that measures beam position directly in units of the beam diameter (see Section 3.5.7) and can be built from scratch in a couple of hours. If you carefully center a biggish cell (10 mm or so) in an aluminum slug the same diameter as the lenses you use most, you can pop it in anywhere in your setup, any time. It works especially well with optical Erector sets such as Microbench.

Germanium quad cells are a good choice if you work in the visible and NIR at reasonable power levels (>100 μW). At very low power levels, the shunt resistance of the germanium may cause the sensor signal to become power dependent. Silicon is better then. The main objection against Ge is leakage, and who cares about that when we're forward biasing them anyway?

12.9.12 Use Fringes for Angular Alignment

Another way to get accurate angular alignment between laser beams is to force them to interfere perfectly. If the beam blinks on and off across the whole diameter when you gently press one mirror, it's aligned as well as it's going to get. A viewer helps a lot with this, and to get absolute perfection you can use a big photodiode, a load resistor, and an oscilloscope; adjust for maximum amplitude as you pound gently on the table.

12.10 ALIGNING LASER SYSTEMS

12.10.1 Define an Axis

In most cases, the best thing to do is to define the optical axis with a HeNe laser. Spend a little time getting the laser aligned properly, really horizontal and really square to the table. If you're using commercial mounts, a height of around 150 mm above the table is a good choice; low enough to avoid floppiness, but high enough to get mounts and stages underneath the beam (a common error is to put the beam too close to the table and then have to suspend later elements from above like freeway signs).

Figure 12.7 shows one good way to align the axis: use a jig consisting of a square piece of sheet metal with one side bent down (or a bar attached) to follow the edge of the table, and an optical mount with an iris diaphragm mounted near the opposite side. Slide the jig up and down the table edge, making the beam go through the iris at all points along its path; this guarantees that the beam is parallel to the edge of the table. An even better method is to substitute the quad cell device of Section 12.9.11 for the iris, which improves the accuracy by two orders of magnitude at least, assuming the alignment beam is clean enough.

Surveyors define the horizontal reference with an autocollimation system consisting of a pentaprism and a mercury bath; shining a HeNe down onto the mercury and adjusting it for autocollimation makes a vertical reference, and an accurate pentaprism turns that into a horizontal reference that does not depend on mild wiggles. In the lab we normally care more about being parallel to the table.

12.10.2 Adding Elements

Start from the far end of the system, and add elements one at a time, making sure the front surface reflections from both sides of all elements go back on themselves (this is not subtle). By the time you get back to the laser, you're all aligned.

This procedure may need to be modified if the optical path is extensively folded; clearly you can't use just the sample and laser if the beam path needs to look like a W.

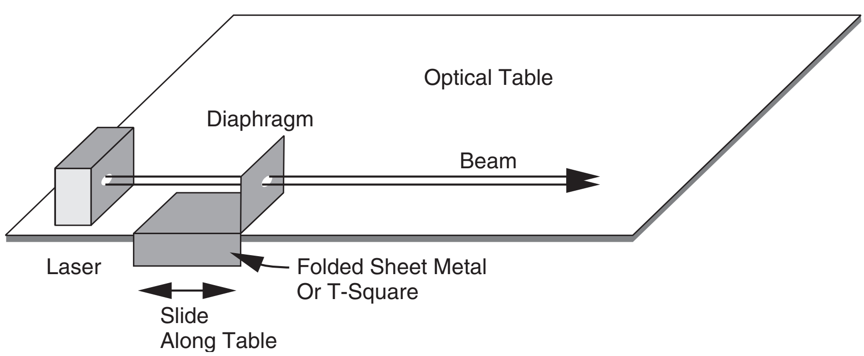

Figure 12.7. Defining an axis.

In that case, set up the folding mirrors and sample, paying close attention to the alignment of each individual section just as before.

This technique works well for beams that are collimated or nearly so; if you have sections where the NA is appreciable (say, 0.05 or greater), build those sections on their own base plates and test them individually.

12.10.3 Marking Lens Elements

You can mark the center of a lens with a small dot of India ink or other easily removed, easily visible material. This helps a good deal with getting the mounts adjusted, or with aligning a previously adjusted subassembly with the system axis. The mark can be positioned by autocollimation testing, or with a lathe or other jig if the lens is sufficiently accurately edged, that is, the axis of the lens coincides with the center of its outline.

12.10.4 Lenses Are Easier than Mirrors, Especially Off-axis Aspheres

Lenses have the nice property that if you shine a laser beam down the axis, the transmitted and reflected rays are at exactly 180° to each other, and the front and back surface reflections overlap completely. Shifting the lens makes the two reflections go opposite ways, and tipping it makes them go the same way. Thus a lens gives you independent cues for position and angle. According to Kingslake,[†] "surface tilt does more damage to an image than any other manufacturing error, and in assembling a lens it is essential to avoid tilting the surfaces at any cost." He suggests limiting accidental surface tilt to less than 1 arc minute, which is easily done using the front and rear surface reflections.

Mirrors have only one surface, so you get only one cue. That isn't too serious with a spherical mirror, of course, because (apart from a bit of vignetting) all points are equivalent; translations and tilts have the same optical effect, so just shine a laser beam down the desired axis and adjust for exact back-reflection.

Aspheric mirrors are much less forgiving in this regard, especially fast ones. It is usually necessary to mark the center of such a mirror somehow, and this is usually done during fabrication. Failing that, you'll probably have to set up a test bench to look at the interference fringes, align it for best performance, and mark it yourself.

Off-axis aspheres are the worst of the lot; since the axis generally lies outside the element, you can't mark it and can't bounce a laser beam off the axial position either. If you're planning to use one of these, make sure you also plan how you're going to align it—you'll probably wind up using the rest of the optical system plus some auxiliary apparatus.

12.10.5 Use an Oscilloscope

Extremely fine alignment (e.g., for an interferometer system) is best done with an electronic sensor of some sort (e.g., a quad cell or a lateral effect cell) and an analog oscilloscope in X-Y mode.[‡] Stick the sensor in the beam, and adjust the beam pointing until the oscilloscope dot is centered in the screen; then jack up the vertical gains and do it again. Obviously you have to use a sensor where 0.000 volts means the light is dead

[†]Rudolf Kingslake, *Lens Design Fundamentals*. Academic, Orlando, 1978, p. 192.

[‡]Turn the display intensity way down so as not to burn the phosphor.

center, and adjust the offsets on the scope so that 0.000 volt on x and y is in the center of the screen. Doing it this way is much easier on the eyes, and it works very repeatably.

12.11 ADHESIVES

12.11.1 Structural Adhesives

Adhesives are very convenient. Need to hold something in an odd position? A drop of Krazy Glue or Duco Cement or a bit of double-sticky tape and you're on your way, right? Actually, adhesives aren't as useful for optical setups as one might think. Glue is dimensionally unstable, first of all. It shrinks a percent or two on curing and keeps creeping slowly with time. It conducts heat very slowly and may delaminate under wide temperature excursions. It also requires extremely careful surface prep for best results—you have to remove oil, fingerprints, and dust, plus roughening the surface everywhere. Scraping it with a scalpel won't do it; the surface area increase is too small. Sandblasting or bead blasting is good, sanding second best. Another problem is outgassing. There exist low-outgassing RTVs such as GE 566 or GE 142; cyanoacrylate cements such as Permabond 910 or Krazy Glue outgas enough to leave white frost on nearby surfaces. Manufacturers laughingly advertise some types as "low outgassing," but that's true only by comparison. Two-part epoxy is usually best, with an hour's baking at 100 °C to toughen it up.

12.11.2 Optical Adhesives: UV Epoxy

One way to simplify the mounting problem, and cut etalon fringes to boot, is to cement your optical elements together. For example, you can carefully cement a $\lambda/4$ plate to a polarizing prism to make a transmit/receive beam separator of the kind we saw in Section 6.10.9.

Find a UV epoxy you like with an index of 1.52 (e.g., Norland 65). This greatly reduces reflections from gluing two pieces of BK7 together, which you'll do a lot. If you're using other glasses, pick a glue whose index is midway between them. It is not unknown for the very weak Fresnel reflections from an epoxy joint to cause etalon fringes bad enough to wreck a spectroscopic measurement, so don't use glued joints the way a witch doctor uses a fetish: calculate the worst case reflection over temperature and design around it.

UV epoxy has only a few months' shelf life at room temperature, but most kinds last for some years in the fridge. Get a good quality two-bulb long-wave UV lamp (black light). If you get one of the enclosed ones with windows, you can rest small parts on the upturned lamp, which is very convenient. Really small parts (<15 mm diameter) can be cured in an EPROM eraser. If you're going to be doing it much, it's worth getting one of the nice little hand-held UV guns Norland sells for tacking the surfaces together, because otherwise they will tend to slide around while they're curing.

UV epoxy is easy to use if you're neat with your hands. Use powder-free finger cots or disposable vinyl gloves while cleaning surfaces and applying glue. The glue is best applied by loading it into a disposable syringe with a filter attached. The epoxy flows freely through filters smaller than 0.5 micron, even down to 0.025 micron, so you can ensure that there are no significant bits of crud in your glue (the crud is usually bits of partly cured epoxy from the spout of the bottle). Get the surfaces good and clean. Using

Opticlean works well, since you can dry-fit the surfaces once they're coated with the polymer film, handle them with your fingers, and so on.

Put one element facing upwards in a soft jawed vice, or wax it to a stable surface. Peel the film from the fixed element, and put on a drop of UV epoxy. If the surfaces are flat or nearly so, put it near one edge. Touch the corner of the top element to the outer edge of the glue, and work it all along one rim. Slowly tilt the top element down so that a bead of glue travels outward in the narrowing wedge-shaped gap, leaving a perfect glue layer behind. (The glue should be applied near the center of the lower element if it's steeply curved, and worked outward with a gentle circular rocking motion.)

Using your alignment jig, slide it around until it's in the right place. Press gently to squeeze the glue layer down, and remove any large globs of excess glue. (The glue layer is called the *bond line* in the trade.) Tack the surfaces together before removing the assembly from the jig, using a portable UV lamp or gun. Tacking is not a full cure but turns some of the glue into a gel stiff enough to hold the pieces together. Epoxy can be removed fairly easily at that stage, so double check the alignment before final cure.

How much glue to use is a matter of judgment. The glue layer should be thin, no more than a mil (25 microns), but you really want to avoid bubbles. It may take a couple of tries to get it right—if you get even one bubble, take it apart, add glue, and try again. It's worth practicing this before hand with corn syrup.

If your assembly needs fine alignment, you should make a gluing jig that will accommodate the alignment procedure you plan to use. After the fine alignment (if any), cure the glue as its maker recommends, and perhaps a bit longer just for luck, especially if your glass is thick or doesn't transmit well at 300 nm. UV epoxy isn't quite as strong as two-part epoxy, but this can be fixed by baking it; 15 minutes to an hour at 100 °C does wonders for the strength of the joint.

Epoxy can actually pull hunks of glass out of the sides of your optics if you leave globs there when you cure it. Make sure any significant amount of excess glue is removed before final curing. Temperature swings can cause delamination of larger optics.

12.11.3 Hydroxyl Bonding

Glass and fused quartz can be bonded very simply with a 1:500 solution of KOH in water, applying it like cement and squeezing the bond line down to <10 nm thick. The OH breaks the surface reconstruction of the silicate, leaving dangling silicate bonds that re-form across the bond line as the solvent evaporates. The resulting bond is covalent, highly transparent, and very strong, but the surfaces have to match very closely to begin with.

12.11.4 Temporary Joints: Index Oil and Wax

For temporary use, you may prefer to use index oil or wax, which are available from Cargille with n from about 1.35 to 1.9 in spacings of 0.002, with waxes reaching 2.0. The low index stuff is pretty safe to use, but some of the higher index ingredients are nastier: read the materials safety data sheet (MSDS) and wash your hands a lot.

A drop of the right index oil will reduce front surface reflections by four orders of magnitude, which is not to be sniffed at. Index oil can also be used to bandage damaged surfaces: scatter from a scratch on a lens can be reduced temporarily by rubbing a very small amount of index oil on it, then polishing it with lens paper,[†] and a pitted window can be helped by putting on a drop of index oil and covering it with a microscope slide.

Oil is applied the same way as UV epoxy, but with still more attention to getting the layer thin, because the surfaces stick together better that way. Even mounted vertically, a good oiled joint with both elements supported from below is stable for months, at least under lab conditions. Wax lasts forever in the lab. You can get much higher viscosity oil for vertical and upside-down use (it's almost as thick as grease).

The temperature coefficients and dispersion of index matching gunk are much higher than those of glass, but that isn't usually a problem; it forms a thin element whose two radii are equal, and the index discontinuity is small, so by the lensmaker's equation it has zero power. Oiled joints also accommodate thermal expansion well.

12.12 CLEANING

We've all seen photographs and slide shows where a large piece of dust or a fiber in the wrong place has made a distracting shadow in the image. The same can happen in instruments, where the consequences are usually far worse than a distraction. Fortunately, most of the surfaces where dust collects are well out of focus, so the shadows are unobtrusive. There is really no way to clean an optic without wear or damage, but some methods are a lot worse than others. Thus it's worth spending some time talking about cleaning.

12.12.1 What Does a Clean Lens Look Like?

Well, if you've ever looked at a camera lens or microscope catalog, you've seen what a lens is supposed to look like. The AR coatings are working beautifully, because there are no fingerprints (which show up bright on an AR-coated surface). We might not have the beautifully positioned lighting of the photographer, but we don't really need it. To examine an optical element, stand with your back turned to an overhead light source (ceiling fluorescents are usually good enough). Rock the element back and forth in your hand, looking at the reflected image of the light source in both the front and rear surfaces (you'll have to refocus your eyes if it's a lens). It should be uniform in brightness and color. Shiny spots are fingerprint grease, which has to go. Next, orient the element so that it reflects some darkish area, for example, the shadow underneath a lab shelf. Rock it back and forth to look for scattered light. Large dust particles are immediately obvious as localized sources of scatter. A broad diffuse scatter is fine dust, condensation (e.g., plasticizers outgassed from plastic mounts), or surface damage such as weathering or etching.

If the lens is to be used for laser work, you have to use a really bright light source to look for even small amounts of scatter. A high powered laser will often burn dust right into the optical surface, destroying it. Scattered light from particles larger than $\lambda/4$ is brighter in the forward direction (small scattering angles), so look on that side. If you're working in the visible at moderate power, use the laser in your system. (Do be careful—most eye damage from lasers happens during setup.) Otherwise a slide projector or microscope condenser works well.

[†]A common photographic darkroom trick to fix scratches on the back of a negative is to smear on a small amount of nose grease with a finger and wipe with lens paper.

12.12.2 When to Clean

Mildly dirty optics are surprisingly innocuous. The areal fraction of the beam that dust occupies is very small, so most of the light is unaffected; the dirt doesn't hurt you until it starts causing so much stray light that your measurement sensitivity is degraded. Dark-field measurements are obviously more vulnerable to dirt than are bright-field ones. The trade-off is between scatter from the dust and scatter from the long thin scratches it will leave when you rub the surface with lens paper or a cotton swab. Scatter tends to depend on the perimeter of the defect more than its area, so this is not necessarily an improvement.

Films are a different matter. A bit of machine oil left in a lens barrel can cause condensation on the surfaces, which really will screw up an optical system. There is a silver lining here, because grease films can be cleaned off without scrubbing (and hence without significant damage), while dust often can't.

Optics don't have time to get dusty when you're actually working on them; most of the dust collects when you're elsewhere. Use dust shrouds, a clean hood, or just a big plastic box with a hinged side that opens when you're working on the system.

12.12.3 Cleaning Lenses

Glass lenses with hard AR coatings on them can be cleaned pretty vigorously without serious damage, as long as you get the dust off first so that it doesn't gouge the surface. Even so, expect the stray light to get slightly worse each time due to accumulated scratches. Rinsing with Kodak lens cleaner or a surfactant solution such as Triton[†] or Kodak PhotoFlo works well unless the lens is heavily fouled with fingerprints, in which case ultrasonic cleaning in a gentle low-suds detergent such as Alconox works well as a first step. Make sure that the optical surfaces don't contact the supports when using ultrasonic cleaning.

If your facility has ultrapure water (18 MΩ deionized or better), then mixing a very small amount of Triton with such good water will make a cleaning solution that will dry without leaving a residue. Mild scrubbing with good quality lens paper, a microfiber cloth, or Q-tips removes most kinds of crud, but will leave fine scratches.

After washing, rinse with very clean water and blow dry with clean dry air or (ideally) clean nitrogen from a LN_2 tank, which is completely oil-free. Don't use machine shop quality compressed air to dry optics.

For everyday cleaning, e.g. getting fingerprints off camera lenses, the microfiber cloths sold for cleaning eyeglasses are quite good. Unlike natural fiber, they don't contain silica and are therefore much gentler. You still have to avoid grinding dust into the lens surface, however.

Some surfaces cannot tolerate such cavalier treatment or cannot be removed for cleaning. Laser tube Brewster windows and soft or water-sensitive IR materials are examples. You can clean these by using a single sheet of lens paper. Put a drop of electronic grade solvent, such as isopropyl alcohol (IPA), acetone, or trichloroethylene (TCE) on the optical surface, lay the lens paper down flat on it (so that the solvent soaks in and wets the whole surface), and slowly drag the lens paper along the surface without applying pressure, as though you were dragging a tarp full of leaves in the garden. Blow dry and repeat, perhaps following a polar solvent such as IPA with a nonpolar one like TCE.

[†]Triton X-100, made by Fisher Scientific, is a very concentrated surfactant that is very clean.

It's important to use electronic grade solvents, since reagent grade ones are not usually pure enough. (Don't use solvents to clean plastic lenses—if you put isopropyl alcohol on acrylic, for instance, it will craze the surface instantly. Stick with aqueous cleaners.)

Some people like to fold the lens paper over several times and hold it in a degreased hemostat, which is easier. If you do this, wear skin gloves while folding the paper, and do it in mid-air. The great thing to avoid in cleaning delicate optics is fouling the lens paper by touching it with your fingers—treat it like a bandage.

Gotcha: Plasticizer Residue. One really insidious source of contamination in a cleaning procedure is leaching of plasticizers from vinyl gloves and plastic parts. These plasticizers are oily substances mixed into plastics to make them soft; the most common one is dioctyl phthalate, also used as vacuum pump oil—its vapor pressure is so low that it stays around forever. As long as they stay put, plasticizers are wonderful, but on your optics, they're a disaster. If you dip your gloved hand into the cleaning solvent, you've just guaranteed yourself dioctyl phthalate all over all of your parts.

12.12.4 Cleaning Gratings

Diffraction gratings must be kept very clean. Contamination will give rise to scatter, and scatter is the enemy of good spectroscopic measurements. The best way to clean a grating is not to get it dirty; cleaning them without damaging them is difficult. A heavily soiled grating may be beyond help even if its surface is otherwise undamaged. Whatever you do, don't touch the grating surface—that's a very expensive mistake indeed.

If the coating is aluminum, the grating should not be soaked long in aqueous solvents because the aluminum may pit slightly, rendering the grating useless. Brief exposure is okay, providing the water is pure enough not to leave residue when it dries—remember, not much residue is needed to produce bright scattered light. If a grating becomes dirty, clean it ultrasonically in a very pure organic solvent such as electronic grade acetone (make sure your ultrasonic cleaner is safe with flammable solvents—some aren't). Gently blow it dry with filtered dry nitrogen or really clean compressed air, chasing the droplets off the ruled surface to avoid residue. Really make sure the air is completely clean and oil-free—one shot of shop-grade compressed air may be enough to kill a grating. Avoid Dust-Off and its ilk, because small amounts of extremely cold liquid may come out with the spray and craze the grating surface.

Some solvents may soften the polymer layer underneath the coating, causing the grating to wrinkle at the edges, so ask the manufacturer first.

Alternatively, you can use a powerful surfactant such as Triton in *very* pure water (18 MΩ deionized preferably, 0.1 MΩ minimum). Avoid using ordinary detergents, or anything but the purest water; they leave spots when they dry, almost no matter what you do. Electronic grade solvents are the cleanest but are not cheap.

Oil contamination wicks down the grating grooves, greatly reducing diffraction efficiency and producing interesting abstract wavefronts. Fortunately, it's not as bad as it looks; rinsing in acetone will completely restore the grating.

12.12.5 Opticlean Polymer

One excellent, although slow, way to get smooth optical surfaces very clean is to use Opticlean Polymer,† which is a solution of polyvinyl alcohol plastic in an organic solvent.

†Opticlean is made by Bradford Laboratories.

You put it on with a swab (a clean room Q tip is best, if you have them), wait for it to dry in a horizontal position, and then peel it off by attaching tape to the top and peeling the tape. The resulting surface is impressively dust-free, suitable for optical contacting or even frustrated TIR measurements.

The two things you have to remember with Opticlean are to use enough to make the film strong, and to be very careful with the edges of the optic. If you let it spill over onto a ground glass area, such as the sides of a prism or the edge of a lens, it will stick tenaciously, and lead to ripping of the film and fouling of the edges. If you're doing a faceted device such as a prism, don't let the films on adjacent facets merge, or it won't strip.

On the other hand, do use enough at the edges. If the film at the edge is too thin, it won't detach well either. The solution is to very carefully form a raised fluid bead at the edges of the facet. This happens naturally with convex surfaces due to gravity, but has to be done manually with flat ones (concave ones can be dried upside down).

The right quantity to use varies with the formulation you choose, but something like a teaspoonful (5 mL) is about right for a 50 mm square area. It looks like a lot more when it's wet. Certain types of Opticlean exhibit unstable fluid surfaces, so that thin areas spontaneously thin further, limited by the rate of drying and the surface tension. With these types, you have to make the layer slightly more carefully, since the thinned areas are difficult to strip off.

Bradford Labs sells its own swabs and tape strips, but they're very expensive. Use a tape that doesn't shed particles: 3M type 471 white PVC tape is very clean (we use it in Class 1 semiconductor clean rooms) and adequately sticky, although stickier would be better. Ordinary wood-handled Q-Tips work pretty well for spreading the polymer, and if they shed a fiber or two, it gets removed with all the rest when you strip the polymer off.

Resist the temptation to try stripping the film before it's completely dry. You can speed it up a little by mild heating (40 °C or less), for example, by pulling your desk lamp down near it. Don't try to force dry it with a heat gun or oven, because it bubbles, spits, and then refuses to come off.

12.13 ENVIRONMENTAL CONSIDERATIONS

12.13.1 Fungus

Most optical systems will develop fungus growth on their surfaces when used in warm and humid conditions. It was worse back when optical cement was made of tree sap (Canada balsam)—in the Pacific battles of WWII, soldiers' field glasses could become unusable in a few days, as fungus ate the cement. It is still common to see camera lenses with spidery fungus growths inside them. Left unchecked, this will eventually etch the glass, destroying the lens. In general, the same measures you take to prevent mildew in cloth appear to help control fungus—keep the optics dry and either (a) sealed or (b) very well ventilated.

12.13.2 Coating Drift

Coatings, particularly interference coatings, exhibit drift with time. They are rarely fully dense, and since they are deposited at high temperatures and low pressures, dielectric coatings are usually somewhat nonstoichiometric and so only metastable in ambient conditions. Over time, these coatings hydrate, causing their indices to go up slightly, and sometimes haze develops. Interference filters that are critical for the long-term performance of your instrument should be protected against humidity if possible. (See Section 5.4.6 for more details.)

12.13.3 Lens Staining

Optical glass is susceptible to staining and corrosion due to environmental exposure. Quartz and crown glass don't stain easily, but flints and filter glass are quite vulnerable—a surface layer forms that gradually becomes opaque. Catalog glasses are rated for stain resistance (see Section 4.2.1).

12.13.4 Drift from Temperature and Humidity

In building a complicated optical system needing fine alignment, it's worth calculating what effects temperature and humidity will have on your alignment. For example, the author has a laser built on a cast aluminum optical bench with a coat of paint on it—under the optical mounts. The paint is relatively thick, at least 125 microns; since the plastic binder can easily expand by several tenths of a percent due to humidity, this is a possible source of instability. (It's a drop in the bucket for this laser, as it turns out, but it might not be for yours.)