CHAPTER 5

Coatings, Filters, and Surface Finishes

Bash to fit, file to hide, and paint to cover.

—Anonymous

5.1 INTRODUCTION

An optical element is just a chunk of stuff put in the way of a light beam. Nearly all the action happens right at the surface, which means that controlling the strength and path of surface reflected and transmitted waves is most of optics. In this schematic view, Chapter 4 is about controlling the path, and this one is about controlling the strength.

The jumping-off point for the discussion of coatings is a more detailed consideration of the Fresnel formulas of Section 1.2.4. From there, we can develop a simple way of calculating the behavior of an arbitrary plane-parallel coating stack.

Besides lenses and mirrors, optical systems use white surfaces, for diffusion, and black ones, for stray light control.

5.1.1 Refraction and Reflection at an Interface

We saw in Section 1.2.4 that the Fresnel formulas (1.8)–(1.11) predict the amplitude and phase of reflected and transmitted plane waves at planar interfaces. Here we'll go into a bit more detail about their behavior. Figures 5.1 and 5.2 show the magnitude and phase of the reflection coefficients at planar interfaces between several pairs of lossless dielectrics, as a function of incidence angle in the higher index medium. It isn't usually plotted this way, because it makes a complicated picture, but there's some useful physics here.

Below the critical angle θ_C, reflections at interfaces between lossless isotropic dielectrics always have phase angles of 0 or π, so that linear polarization stays linear after encountering such a surface. Above there, the phase goes all over the place, as you can see; the polarization change on TIR depends on the phase difference δ between s and p, so keep careful track of polarization when using TIR prisms. (This isn't all bad; in Section 4.9.10 we use this effect for polarization control with Fresnel rhombs.) Note especially that r_p goes negative between the Brewster angle θ_B and θ_C.
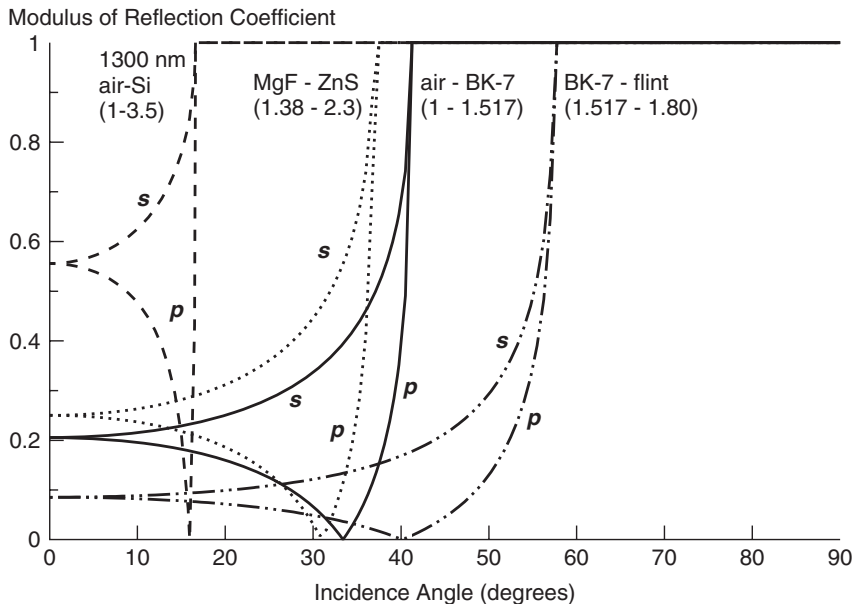
Modulus of Reflection Coefficient

Figure 5.1. Modulus of the reflection coefficients $|r_p|$ and $|r_s|$ at a dielectric boundary, for several choices of material. The light is assumed to propagate from higher n to lower.

Reflection Phase (degrees)

Figure 5.2. Phase of the reflection coefficients r_p and r_s at a dielectric boundary, for several choices of material. The light is assumed to propagate from higher n to lower.
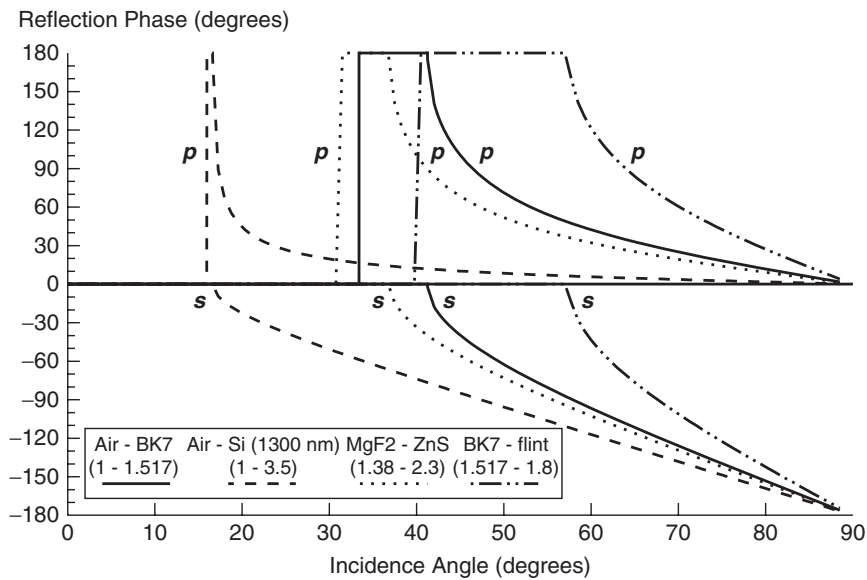
5.2   METAL MIRRORS

5.2.1   Lossy Media

The most familiar mirror coatings are metals such as silver, aluminum, and gold. From an optical point of view, the distinguishing feature of a metal is the very large imaginary part of its refractive index, that is, its extremely high loss. It is slightly paradoxical at first blush, but the highest reflectance coatings are made from the highest loss materials. (See Figure 5.3)
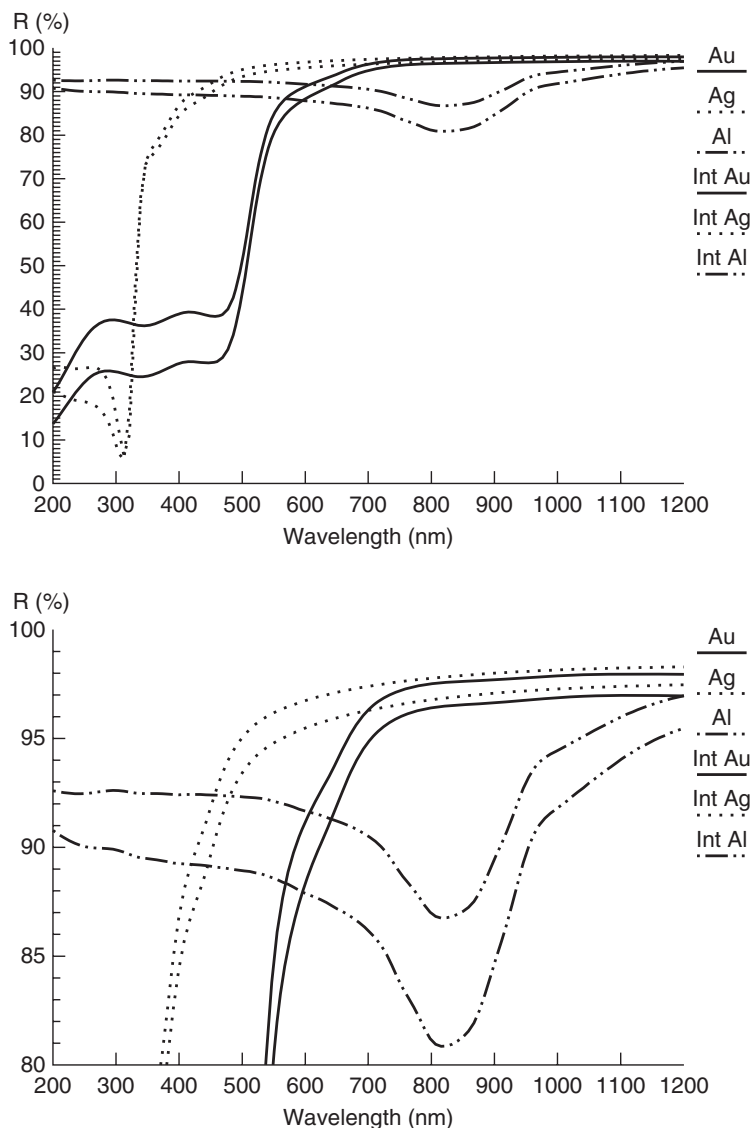
Figure 5.3. Theoretical normal incidence reflectance of gold, silver, and aluminum mirrors as a function of wavelength, for external and internal ($n_{glass} = 1.52$) reflections.

The Fresnel formulas in terms of θ_i are also valid for absorbing media, even metals. Aluminum has an index of about $0.76 + i5.5$[†] for green mercury light (546 nm), so a clean air–aluminum interface at normal incidence has a reflection coefficient of

$$r = \frac{n_2 - n_1}{n_2 + n_1} = \frac{-0.24 + i5.5}{1.76 + i5.5} = 0.953\angle 160°, \tag{5.1}$$

and so the reflectivity $R = |r|^2 = 0.91$. An internal reflection (e.g., an aluminized BK7 glass prism) has an even lower R of 0.87.

From a designer's point of view, just bouncing a light beam once from an aluminum mirror costs 0.8 dB in detected (electrical) signal power for a first surface mirror and 1.2 dB for Al–glass. You can't do that too many times and still have a signal; this is one of the reasons for the popularity of folding prisms based on TIR, especially in complicated optical systems like microscopes. Since n_1 is real, the reflection coefficient is maximized for a given $|n_2|$ when n_2 is purely imaginary. Metals work a whole lot better in the IR.

5.2.2 How Thick Does the Metal Have to Be?

We saw that metals make reasonable although not fabulous mirrors. Since we usually don't want to make mirrors from solid metal, we need to know how thick to make the coating. The transmitted light typically has an initial amplitude $|E''| \approx 0.2|E|$, so we can't just ignore it—we have to make the coating thick enough to be opaque.

Because of the hugeness of $\text{Im}\{n\}$, the wave equation in metals behaves like the diffusion (heat) equation, so that the electromagnetic fields are diffusive in character[‡]; the amplitude of a plane wave in a metal dies off by $\exp(-2\pi)$ per cycle. Also, $|n|$ is so large that **k** is directed essentially normal to the interface, regardless of incidence angle.

In order for the light transmitted through the film not to significantly reduce the reflected power, we choose a film thickness $d > \lambda/\text{Im}\{n\}$, so that the light making a round trip through the film is attenuated by at least $e^{-4\pi}$, or 3×10^{-6}. A thousand angstroms of aluminum makes a good mirror in the visible, but 200 Å is getting a bit see-through. The optical constants of metals change so rapidly with wavelength that the thickness required can easily vary 2:1 across the visible, so light transmitted through a thin metal layer is often strongly colored. Metallic neutral density filters are made of chromium, rhodium, Inconel, or other metals whose properties are less wavelength dependent.

Aside: Free-Electron Metals. In the infrared, metals such as copper, silver, and gold exhibit *free-electron* behavior. That is, their dielectric constants behave as though the electrons were undamped harmonic oscillators. Since the harmonic oscillator equation is of second order, the response to a sinusoidal **E** at a frequency above the resonance is multiplied by two copies of $j\omega$—in other words, these metals have *negative real dielectric constants*. This leads to all sorts of useful effects such as *surface plasmons*, which are true electromagnetic surface waves that may turn out to have lots of useful technological applications.

[†]With our Fourier transform sign convention, absorbing media always have positive $\text{Im}\{n\}$. Why?
[‡]*Diffusive* is used in its mathematical rather than optical sense here.

5.2.3 Designing Metal Films

Silver is the best metal coating in the visible but tarnishes so fast that it is nearly useless for first-surface reflections. It works much better than aluminum ($R = 0.96$) for internal reflectors (e.g., pentaprisms), where the back of the silver can be protected with a plated layer of less active metal (e.g., copper, Inconel, or nickel) and a coat of paint. On the other hand, if silver is protected from sulfides in the air, it lasts for months, and it can be applied chemically in the lab, which is sometimes a very important property.

Gold is the best mirror coating in the NIR and the red ($\lambda > 633$ nm). It is often convenient for lab use, as many laboratories have small gold sputtering chambers intended for electron microscope samples, so the turnaround time is unbeatable. It doesn't stick all that well, though, so be careful with it. A very thin adhesion layer of chromium or titanium makes gold stick to glass very well. Rhodium is a noble metal whose reflectivity holds up well far into the UV (that's the good news—the bad news is that $R \approx 0.8$ in the visible and 0.4 in the UV). Good metal films (e.g., Cu, Ag, Au, Al) in the IR are essentially perfect conductors, so their efficiency is excellent.

5.3 TRANSMISSIVE OPTICAL COATINGS

Most optical surfaces are not bare glass or plastic, but are coated with one or more thin layers of another material to modify their transmission and reflection properties. The most common are antireflection (AR) and mirror coatings, but sometimes we want a beamsplitter, a polarizer, or a filter, all of which can be made by appropriate coatings.[†]

One can treat coatings in different ways; because we're after physical insight rather than, say, efficient simulations, we'll keep it simple and use the Fresnel formulas, assuming a plane-parallel geometry and homogeneous, isotropic films and substrates.

5.3.1 Dielectric Coating Materials

The theoretical performance of coatings is limited by the available materials and by the difficulty of getting good coating morphology and adhesion. In the early days of coated optics (the 1940s), the best available method for putting down dielectric coatings was vapor phase reactions in air. This produced surprisingly good coatings, so good that Carl Zeiss gave up on vacuum coaters completely for a while (see Anders).

At present, most coatings are deposited in a vacuum, by evaporation or sputtering. These techniques work well, but almost all coatings are a bit less dense than the bulk material, owing to the formation of small voids during deposition. These voids reduce the refractive index slightly and, by adsorbing air and water, cause drift in n with temperature and humidity. They also reduce the corrosion protection the coating affords. This is not especially serious with substrates that form thin oxides with good properties (e.g., Al), but is more of a problem with silver and copper, which do not.

Films are highly variable, depending on deposition conditions such as stoichiometry, humidity, choice of substrate, temperature, pressure, and other things, and there is significant variation from run to run. The microstructure of the film may be quite different from the bulk material's; a good quality coating is amorphous, which will influence its

[†]The discussion of optical coatings is indebted to the admirable small monograph by Hugo Anders of Carl Zeiss, Oberkochen, *Thin Films in Optics*, Focal Press, London, 1967 (J. N. Davidson, tr.).

refractive indices and transmission bands significantly (especially for anisotropic materials). Besides porosity and microstructure, coatings often have stoichiometry errors that can make *n* go up or down. For example, one maker quotes a range of 1.34–1.38 for its MgF_2 coatings at 550 nm. The idea of a table of optical constants of thin films is thus something of an oxymoron, so don't take the values in Table 5.1 too seriously. Above all, don't expect to get exactly the same refractive index and transparency range as the bulk material. The high variability of the indices of the films, and the moderate

TABLE 5.1. Common Coating Materials

Material	Index	λ (nm)	Comments
Cryolite (Na_3AlF_6)	1.35	Visible	Lowest *n* among dense coatings; water soluble; soft
Magnesium fluoride (MgF_2)	1.38	Visible	Lowest index hard coating; popular
Quartz (SiO_2)	1.46	Visible	
Silicon monoxide (SiO)	1.5–1.9	Visible	Nonstoichiometric; high index SiO absorbs in the blue
Silicon nitride	2.02	500	Absorbs strongly below 300 nm
Sapphire (Al_2O_3)	1.75	Visible	
Titanium dioxide (TiO_2)	2.5	Visible	Varies between 2.2 and 2.7
Zinc sulfide (ZnS)	2.35	Visible	
Lead fluoride	1.75	Visible	
Indium–tin oxide	2.0	500	Electrically conductive; good ITO transmits >87% in the visible
Silicon (Si)	3.5	1300	
Gold (Au)	$1.66 + i1.956$	400	Needs Cr or Ni adhesion layer; best metal for λ > 900 nm
	$1.24 + i1.80$	477	
	$0.61 + i2.12$	517	
	$0.31 + i2.88$	564	
	$0.16 + i3.80$	689	
	$0.19 + i5.39$	827	
	$0.27 + i7.07$	1030	
	$7.4 + i53.4$	10,000	
Silver (Ag)	$1.32 + i0.65$	310	Corrodes rapidly; best metal in the visible; much poorer below 400 nm; can be applied chemically in the lab
	$0.17 + i1.95$	400	
	$0.13 + i2.72$	476	
	$0.12 + i3.45$	564	
	$0.15 + i4.74$	729	
	$0.23 + i6.9$	1030	
	$0.65 + i12.2$	2000	
	$10.7 + i69$	10,000	
Aluminum (Al)	$0.13 + i2.39$	207	Reasonably stable in dry air; best all-round metal; reflectance dips badly in the deep red and near IR (700–1000nm)
	$0.49 + i4.86$	400	
	$0.76 + i5.5$	546	
	$1.83 + i8.31$	700	
	$2.80 + i8.45$	800	
	$2.06 + i8.30$	900	
	$1.13 + i11.2$	1130	
	$25.4 + i67.3$	10,000	

difficulty of obtaining highly uniform films of the correct thickness, must influence the way we design with thin films—a design needing three- figure accuracy in refractive index and thickness will be totally unmanufacturable. Theoretical elegance must be sacrificed to the exigencies of coating manufacturing. (This is not a blanket dismissal of fancy coatings—some highly multilayer coatings have been designed precisely to be very tolerant of certain classes of coating errors.)

Beyond the electromagnetic properties of coatings, their mechanical ones, such as adhesion, residual stress, and environmental sensitivity, must be considered. Not every coating material sticks to every other, or to every substrate (a couple of nanometers of Cr, Ti, or Ti_3N_4 can help a lot with metals). Materials with a coefficient of thermal expansion (CTE) very different from that of our glass will experience severe stress upon cooling, which may make the film craze or delaminate. Surface preparation is vitally important too—coating really is a bit of a black art. Detailed design of thin film coatings is beyond our scope, because most instrument builders buy parts already coated; nevertheless, you may easily require a custom coating occasionally and so need some understanding of the difficulties involved.

5.4 SIMPLE COATING THEORY

Grinding through the algebra for multiple-layer coatings becomes less fun very rapidly. What's more, the formulas so obtained provide no insight and are useless in practice due to their specificity. You can formulate the problem as a band-diagonal matrix equation based on the matching of tangential **E** and perpendicular **D** at each interface, with a phase delay of $\exp(\pm i k_z z)$ for waves going in the positive and negative z direction, but there's an easier way: for a plane-parallel geometry, the Fresnel formulas can be cobbled together to give us the full electromagnetic solution, assuming that all the materials are isotropic. We'll add up all the reflections and find the result, taking account of the propagation phase delays. As usual, we take the z direction to be the surface normal directed from layer j to layer $j + 1$.

A wave $\exp(i\mathbf{k} \cdot \mathbf{x} - \omega t)$ delayed by propagating through a layer of index n_j and thickness d_j acquires a phase delay of $k_{Zj}d_j$. The value of k_{Zj} depends on n_j, so we have to find it using the phase matching condition.[†] Phase matching states that $\mathbf{k}_\perp$ is preserved across a boundary, as is necessary to preserve translational invariance. Thus in the jth layer, k_{Zj} obeys

$$k_Z^2 = n_j^2 k_0^2 - |\mathbf{k}_\perp|^2. \tag{5.2}$$

All the forward waves in the jth layer have this k_Z, and the reverse waves (i.e., those reflected an odd number of times) have $k'_{Zj} = -k_{Zj}$.

A point of terminology: film thicknesses are universally specified in waves and not nanometers: a half-wave film is one whose optical thickness is $\frac{1}{2}$ wave, that is, $d_j = \lambda/(2n_j)$. In coatings such as beamsplitters, intended to be used off-normal, a "half-wave" coating is usually specified as $d_j = \pi/k_{Zj}$, that is, thicker by a factor of sec θ_j, where θ_j is the angle of incidence in medium j.

[†]We're doing wave propagation here, so we'll stick with the physicists' sign convention, where a plane wave is $\exp[i(\mathbf{kx} - \omega t)]$, and a time delay τ contributes a factor of $\exp(+i\omega\tau)$.

5.4.1 Multilayer Coating Theory

Multilayer coatings appear much more complicated, and in fact the explicit formulas for their reflection and transmission coefficients are ugly enough to give small children nightmares.[†] We therefore proceed recursively, calculating the effect on an existing coating stack of adding another layer, using the total reflection coefficient $\hat{r}$ of all lower layers in place of the simple r from the Fresnel formula on the bottom side (see Section 1.2.4).

First we need to restate the Fresnel formulas ((1.8)–(1.10)) in terms of **k** instead of θ_i; we especially must keep the index *12* for propagating from n_1 into n_2, since as we saw the Fresnel coefficients are not the same going in and coming out, and in dealing with multiple bounces we'll need to keep them straight:

$$r_{p12} = \frac{n_2^2 k_{Z1} - n_1^2 k_{Z2}}{n_2^2 k_{Z1} + n_1^2 k_{Z2}}, \qquad r_{s12} = \frac{k_{Z2} - k_{Z1}}{k_{Z2} + k_{Z1}},$$

$$t_{s12} = \frac{2k_{Z1}}{k_{Z1} + k_{Z2}}, \qquad t_{p12} = \frac{2n_1 n_2 k_{Z1}}{n_2 k_{Z1} + n_1 k_{Z2}}. \tag{5.3}$$

Referring to Figure 5.4, we observe that the multiple bounces form a geometric series with common ratio $\hat{r}_{23} r_{21} \exp(+i2k_{Z2}d_2)$, so that the total reflection coefficient $\hat{r}_{12}$ can be expressed as the total reflection coefficient $\hat{r}_{23}$ of the lower system with the additional effect of the layer n_2, yielding

$$\hat{r}_{12} = r_{12} + \frac{t_{12} t_{21} \hat{r}_{23}}{\exp(-i2k_{Z2}d_2) - \hat{r}_{23} r_{21}}. \tag{5.4}$$
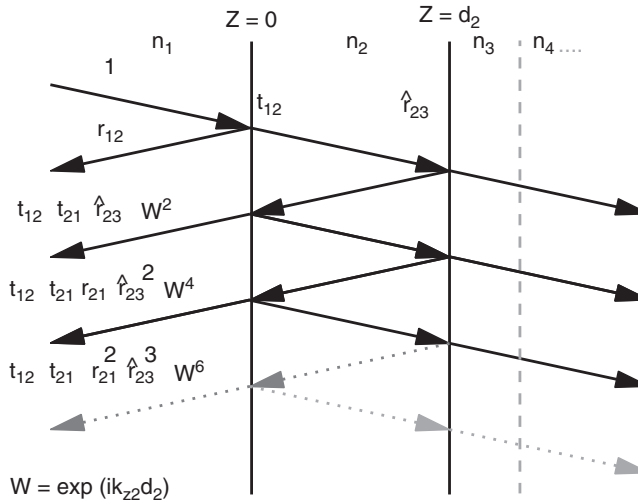

Figure 5.4. Geometry of the plane-parallel coating problem, showing the geometric progression of multiple bounces from the stack below.

[†] Besides, the amount of CPU time spent calculating them is minuscule compared with the time spent getting the coating recipe debugged.

For the bottom layer, we take $\hat{r}_{12} = r_{12}$ from the Fresnel formula, and for all subsequent ones we use the total reflection of all lower layers, obtained from repeated application of (5.4), which gives us a recursive algorithm for calculating the total $\hat{r}$ for the stack. We can compute t the same way, but now we need to go from the top of the coating stack toward the bottom, instead.

5.4.2 Lossless Coating Examples

Here are some typical examples of the uses of coatings. The goal is physical insight, not detailed coating recipes, so we neglect dispersion, material absorption, adhesion problems, and interface effects, for example, the 10–50 nm of Al_2O_3 that grows immediately on top of deposited aluminum, even in vacuum. Don't take this to mean that these effects are negligible in practice.

Example 5.1: Single-Layer AR Coating. For a single interface, the strength of the Fresnel reflection depends on n_2/n_1 and θ_i. At normal incidence, we can get reflections of the same strength and phase from both sides of a coating if $n_3/n_2 = n_2/n_1$, that is, when n_2 is $(n_1 n_3)^{1/2}$. By choosing the layer to be $\lambda/(4n_2)$ thick, the two reflected waves are put π out of phase (there and back) and thus cancel perfectly. The problem is that low index glass requires that n_2 be around 1.2 or 1.25, and there aren't any solid materials in that range (people have used silica aerogels with some success in the red and IR). The lowest index material that is hard and insoluble is MgF_2, at 1.38. This is an excellent AR coating for high index glass, but it's no great shakes with garden-variety borosilicate such as BK7 (1.517), as you can see from Figure 5.5 (We'll take 1.37 as its index, a reasonable value closer to the center of the distribution than the bulk value of 1.38.) Note the angle tuning; at normal incidence the coating is centered around 514.5 nm, but at larger angles it shifts significantly to the blue. Higher index materials have more constant k_Z (since $k_\perp \leq k_0$). The coating is also polarization sensitive, which means that
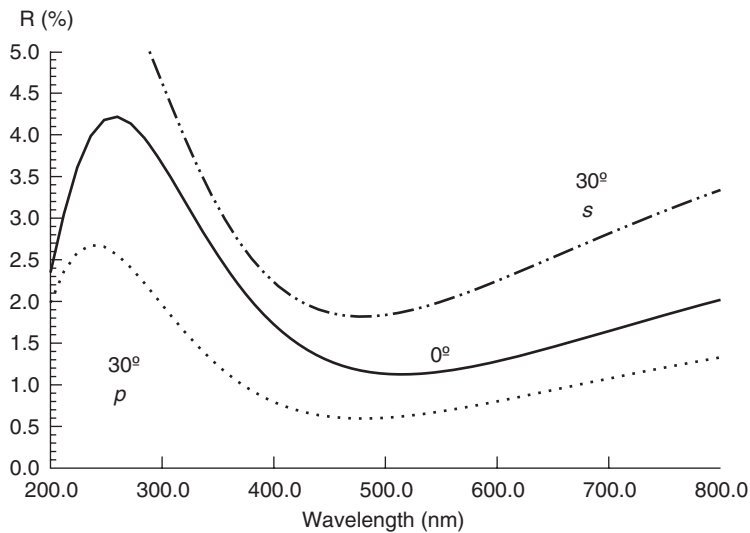

Figure 5.5. Single layer MgF_2 coating on BK7. Note the angle tuning and polarization dependence.

the polarization of your incoming beam will be changed somewhat by passing through the coated surface. A number of such surfaces together (e.g., in a camera lens) can easily cause serious polarization shifts with position and angle.

Example 5.2: Protected Aluminum Mirrors.   First surface metal mirrors are too soft to clean vigorously and (except for gold) are also vulnerable to corrosion. The usual solution is to put a dielectric overcoat on top for protection. We saw that a glass–metal interface was not as good a mirror as air–metal, so protected metal mirrors start out with poorer performance even than Al–air. We can adjust the relative phases of the two reflected waves by changing the thickness of the coating; if we make the layer $\lambda/2$ thick, the reflections will add in phase. This effect is used to make hybrid mirrors, where the coating partly offsets the poorer reflectance of an internal reflection from aluminum. The most common coating choice is $\lambda/2$ of SiO over Al, the classical *protected aluminum* coating of Figure 5.6. Over a relatively narrow bandwidth, the result can be as good as a bare aluminum coating (the "internal reflection" curve is calculated for Al/SiO). These mirrors are OK for simple systems, or ones in which you have photons to burn; be careful how many bounces you use, though, or you may burn more than you can spare. The coating is cheap, which is important, but that and physical robustness about exhaust its virtues. For polarization-sensitive systems such as spectrometers, the fact that a significant proportion of the reflectance comes from the thin film coating means that protected aluminum mirrors are partially polarizing when used off normal. This is not so in the IR, where metals are more or less perfectly conducting; there, a dielectric coating of any thickness does not change R, which is always 1; protected gold coatings make great IR mirrors.

5.4.3 Angle Tuning

Because of the variation in k_Z with incidence angle, the tuning of coatings shifts to shorter λ as θ_i increases, a phenomenon called *angle tuning*. It's an easily calculated effect that can be reduced by using high index materials and reducing the field angle.
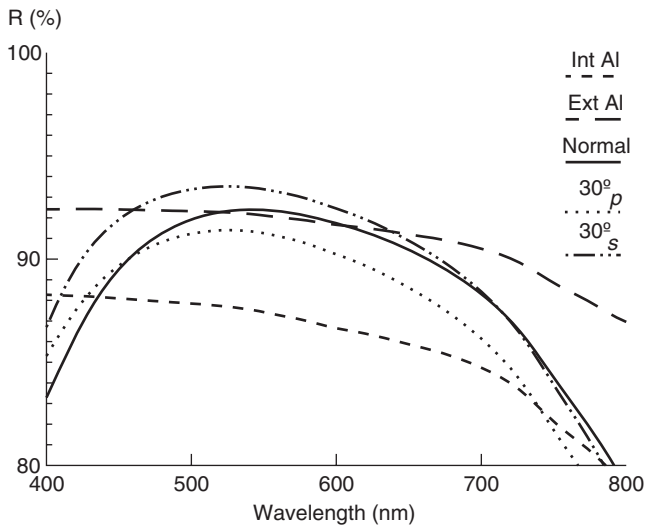

Figure 5.6. The protected aluminum mirror: 0.5 wave at 520 nm of SiO ($n = 1.7$) over Al.

Tuning with angle is generally a minor nuisance, as in Example 5.1, because k_Z is a weak function of $k_\perp$ at small angles. The difference between r_p and r_s and the increase in $|r|$ at high angles usually cause us worse problems. There are exceptions to this rule, such as polarizing beamsplitters, in which selectivity degrades with angle, and sharp interference filters, which can angle-tune your signal right out of the passband.

5.4.4 Examples of Multilayer Coatings

A good coater can put down layers of different materials without breaking vacuum, so it is convenient as well as worthwhile to put down stacks of many layers. Excellent antireflection coatings require several layers, and more complicated multilayer coatings can have excellent performance as mirrors, filters, and beamsplitters. A good AR coating can achieve 0.2% reflectance over a 50% bandwidth, with 0.4% being a typical guaranteed spec. Coatings for use across the whole visible usually achieve <1%.

Most optics manufacturers have a variety of standard coatings available, which may or may not be stock items. The price list will specify how much the coated elements cost, but be aware that lead times are frequently rather long (several weeks is not unusual). If you want different coatings on different surfaces, or a custom coating on parts you supply, be prepared to pay for it: even apart from design time, you'll be charged a setup fee of perhaps $1200 for a coating run, plus a per-piece charge of around $100 per surface, and will get no guarantee that your parts will survive.

Example 5.3: V-Coating. Single-layer AR coatings on plastic and low index glass don't work too well, because there are no materials available with n near 1.25. If we care about a narrow range of wavelengths (e.g., in a laser system), a quarter-wave of MgF_2 over a quarter-wave of SiO (with $n = 1.71$) can fix this problem, as Figure 5.7 shows. Here the coating's reflectance is calculated at normal incidence, and at $30°$ in both s and p. The coating angle-tunes to shorter λ with increasing θ_i, as we expect. Note how the p
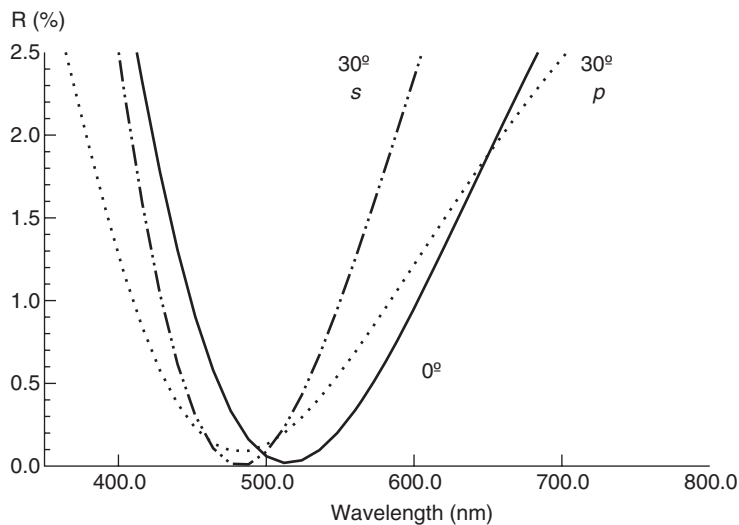
Figure 5.7. Two-layer V-coating: quarter-wave MgF_2 over quarter-wave SiO (1.70) on BK7.
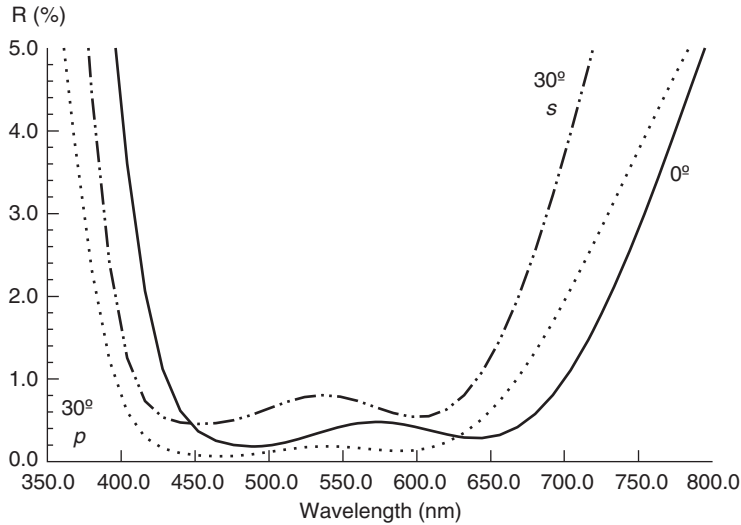
Figure 5.8. BBAR coating.

polarization is more tolerant of angular shifts; the decrease in the r_{pi} values partially compensates for the error in k_Z.

Example 5.4: Simple Broadband AR Coating.   For crown glass (1.46–1.65), a four-layer stack consisting of 0.50λ of SiO (1.60), 0.54λ MgF$_2$, 0.25λ SiO (1.60), and a final 0.27λ of MgF$_2$ on top makes quite a reasonable broadband AR (BBAR) coating; Figure 5.8 shows the calculated results on BK7 (1.517) at normal incidence, and $30°$ *s* and *p*.

Example 5.5: Enhanced Aluminum.   The idea of metal with a dielectric overcoat can be improved by using an *LH* pair over a metal surface: Figure 5.9 shows the result of using a quarter-wave each of ZnS over MgF$_2$ on top of the aluminum (as before, the "internal reflection" curve is calculated for the index of the film adjacent to the aluminum). Another *LH* pair on top produces an even better mirror, with $R > 95\%$ even in the 800 nm dip. Enhanced aluminum mirrors are nearly as cheap as protected aluminum and are the minimum quality you should consider for high sensitivity instruments. (See Figure 5.9.)

Example 5.6: Quarter-Wave $(HL)^m H$ Stack.   A sufficiently tall stack of $\lambda/4$ layers of alternating high and low index materials makes an excellent mirror. It works better with high index layers on both ends of the stack, so that the optical prescription is $(HL)^m H$. Figure 5.10 shows an 11-layer ($m = 5$) stack tuned to 514.5 nm. Note that it's now the *p* polarization that droops at high angles, since the r_p values are dropping as θ_i increases.

Example 5.7: Stagger-Tuned HL Stack.   Although the HL stack high reflector becomes broader as the number of layers increases, this happens only as $\sqrt{N}$, which is wasteful, and even this eventually stops due to absorption in the coating. By putting two HL stacks,
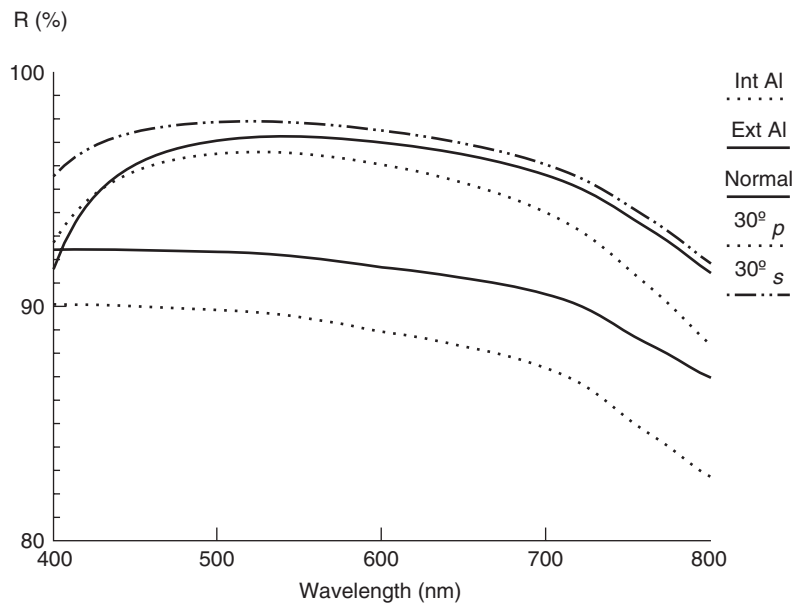
R (%)

Figure 5.9. Replacing the half-wave of SiO with quarter-waves each of ZnS over MgF_2 yields a substantially improved metal mirror for the visible, the enhanced aluminum coating.
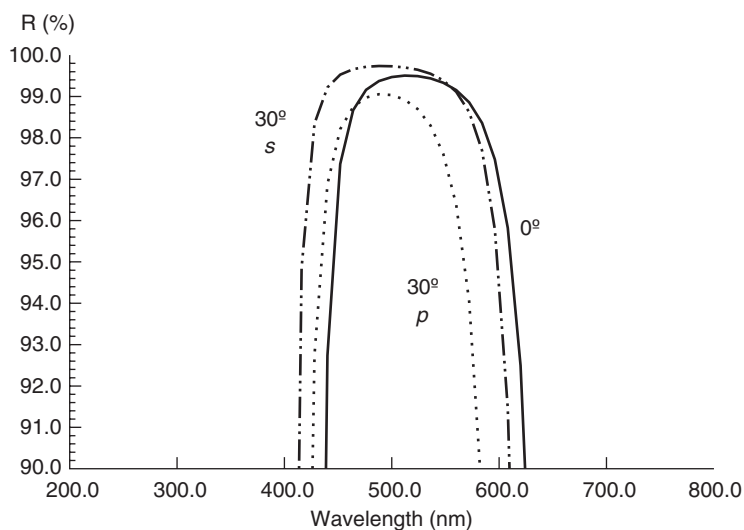

Figure 5.10. Eleven-layer $(HL)^5H$ stack, ZnS/MgF_2, centred at 514.5 nm.

tuned slightly differently, we can get a broader bandwidth with very high efficiency, as we see in Figure 5.11. You usually use a spacer layer between them. This idea is called *stagger tuning*, and it is broadly useful as a bandwidth-increasing device, in circuits as well as optics.
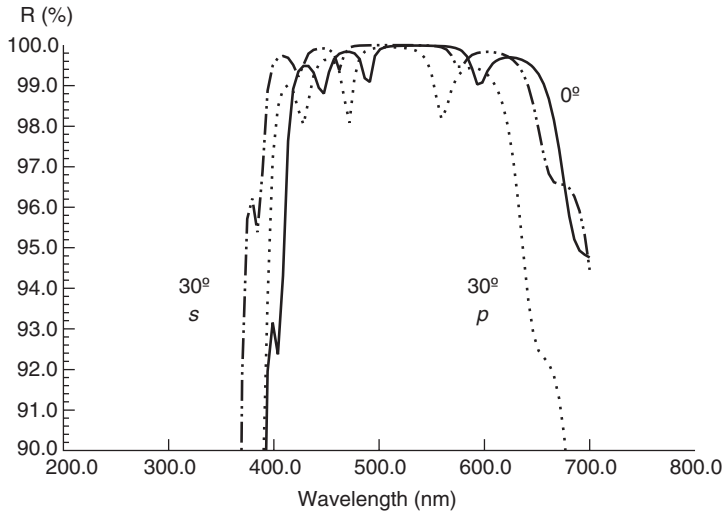
Figure 5.11. Two $(HL)^5 H$ stacks (ZnS/MgF$_2$), tuned to 463 and 600 nm, with an 0.21 wave L spacer layer.

5.4.5 Polarizing Beamsplitters

Beamsplitters used to be made with thin metal films, and cheap ones still are. Inconel is the most popular choice due to its spectral neutrality, but all such coatings are very lossy—typically 40–50% gets absorbed, which is a lot, considering that even a lossless beamsplitter wipes out 75% of our light if we go through and back. Modern beam-splitters are made from dielectric stacks for narrowband or polarizing applications (see Figure 5.12), and from even thinner metal films with dielectric overcoats for wideband applications. You typically lose 15–25% of your light in an enhanced metal beamsplitter, and 5% or less in a good polarizing one.

Because r_s and r_p are so different at high incidence angles, the performance of coatings changes drastically with polarization at high angles. This is how broadband polarizing beamsplitter cubes are made: you come in at 45° to an $(HL)^m H$ stack, with enough layers that T_s is very low[†] (11 layers ($m = 5$) of ZnS and cryolite will get you to 0.0001, not counting absorption), for example, the one in Figure 5.11.

You choose n_{glass} so that light coming in at 45° to the hypotenuse (i.e., normal to the cube face) is refracted into the coating stack at Brewster's angle for the HL interface; this guarantees Brewster incidence at lower layers, because there are only two different indices involved. Cementing another 45° prism to the top of the coating stack makes a cube, where the transmitted light is undeviated and the reflected light comes off at 90°. This makes a good broadband beamsplitter, whose main problem is the first-surface reflections at the glass–coating and coating–cement interfaces. These pollute the reflected light with as much as 5% of the p polarization (it varies with λ because the two p reflections interfere). Adding extra AR coatings top and bottom can make a very good beamsplitter,[‡] marred only by its narrow angular acceptance (and, of course, the horrendous etalon

---

[†]Remember that you have to rejigger the coating thicknesses a bit to make $d_j k_{Zj}$ equal to $\pi/4$ at each layer.
[‡]Not even close to a Wollaston prism for p-polarization quality, of course, but lower in absorption and cheaper—pretty good for a chunk of glass with a few films on it.
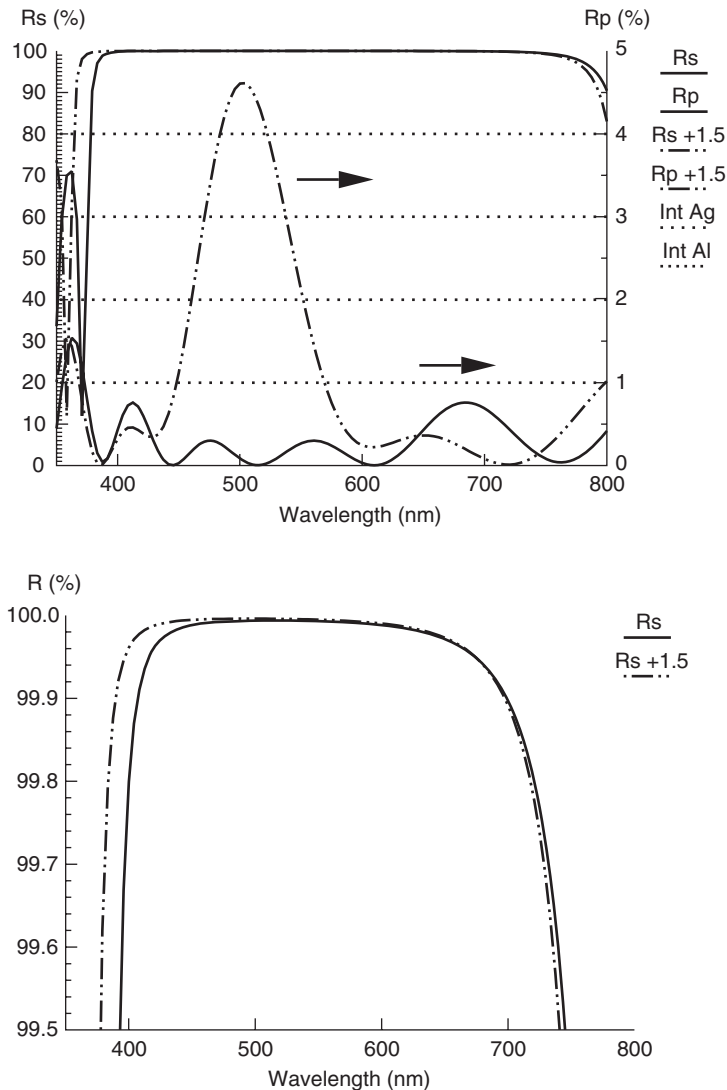
Figure 5.12. Polarizing beamsplitter: $A(HL)^5HA$, where A is SiN (2.0), and H and L are ZnS and cryolite. The AR layer A suppresses the reflection from the glass (1.66). Note the selectivity reduction (to 20:1 from 100:1) due to coming in at only $1.5°$ off normal.

fringes due to the manufacturer's insisting on choosing $45°$, which we've alluded to in Section 4.7.2).

You can also make wider angle, narrower band beamsplitters by working off Brewster's angle, at a wavelength where r_p has fallen way off but r_s is still large. By careful control of the sidelobes of the $(HL)^mH$ stack's reflectance, you can make good beamsplitters for laser applications this way.

The polarization purity is still worse in the reflected light, only 25 or 50:1, whereas in the transmitted light it can be 1000:1 in a narrowband device or 100:1 in a wideband one.

Aside: Unintentional Polarizing Beamsplitters.   Some coatings are pretty good polarizers, including yours if you're not careful. Polarization effects in AR coatings cause major problems with lasers and in high accuracy applications.

5.4.6 Interference Filters

Two HL stacks separated by a spacer layer make a Fabry–Perot etalon, which has a sharply peaked passband near where the spacer is an integral number of half-wavelengths thick (phase shifts from the *HL* stack move the passbands around a bit). The bandwidth and free spectral range can be traded off by changing the spacer thickness, from $\lambda/2$ up.

A complex structure composed of two of these etalons deposited on top of each other, so that their passbands coincide at only one peak, makes an excellent filter with adjustable parameters. As usual, stagger tuning can flatten out the peak and suppress the sidelobes, yielding a flat-topped bandpass filter with steep edges and few artifacts.

Interference filters are fiddly to develop and need long coating runs that must be precisely controlled. Thus they tend to be available only for commonly desired wavelengths (e.g., laser lines) and spectral bands of wide interest (e.g., Balmer α at 656 nm). They are normally backed by colored glass, to suppress unwanted sidelobes, so that out-of-band light hitting the top of the filter is mainly reflected, while that hitting the bottom is largely absorbed. Sometimes it matters which way round you put the filter, for example, in the infrared, where the absorbing glass radiates but the mirror coating doesn't, and with high powered light sources, which may overheat the filter.

The stopband rejection of interference filters isn't always that great. Besides the occasional spurious peak, they sometimes have only 30 dB (optical) typical rejection. That might not be too bad in a color compensating filter for a camera, where the passband is wide and the rejection requirements modest. On the other hand, if you're looking at solar $H\alpha$ absorption at 656 nm with a 0.5 nm passband, you're in trouble—30 dB rejection means that each nanometer in the stopband will be attenuated by 30 dB, but there are a lot more nanometers in the stopband than the passband, so the background light will dominate. Make sure you calculate and measure the total out-of-band leakage in your filters. A quick test is to cant the filter enough to angle-tune your desired signal into the stopband, and see how much the signal level changes. This isn't really precise, because the *L* layers angle-tune more than the *H*, so the shape of the curve will change with angle too.

If you use interference filters in your instrument, be aware that they drift with temperature and time. They normally drift toward longer λ with increasing *T*, generally with $\Delta\lambda/\lambda \approx 10\text{–}30$ ppm/K; time variations are usually associated with hydration (or even corrosion) of the coatings (see Section 12.13.2). Get detailed specs from the manufacturer.

5.4.7 Coating Problems

Coatings usually have a columnar morphology, which makes them porous, chemically somewhat unstable, nonstoichiometric, and often having properties significantly different from the bulk material. Lots of work has gone into fixing these problems, but the solutions are different for different coatings. Sometimes depositing at an angle of $30°$ or so, or using ion bombardment during deposition, can reduce porosity and produce a fully dense coating. The columnar morphology can be exploited (e.g., by rotating the substrate eccentrically at an angle to the deposition source), so as to make the columns helical—that makes an optically active coating (see Section 6.3.6).

5.5 ABSORPTIVE FILTERS

Optical filters are used to get rid of light we don't want. (They always get rid of some of the desired light as well, but not too much with a bit of luck.) Here we'll talk about absorbing materials and scattering from small particles.

5.5.1 Filter Glass

Glass makers can exploit the absorption characteristics of different materials to make intentionally colored glass filters. Filter glass comes in a wide range of colors and characteristics, but the two most used are long pass and short pass, with the long pass glasses having generally better performance. The coloring can come from one of two sources: colloids, which are formed by a carefully controlled heat treatment process (*struck* or *colloidally colored* glass), or by the formation of color centers due to ionic doping (*ionically colored*). Ionically colored glass is a great deal more stable with time and thermal history. Color centers are not easily bleached by optical dose either, so ionically colored glass is pretty stable all round. In addition, it can be annealed to eliminate stress birefringence.

The data sheet for the filter glass usually tells how the color has been achieved. The transmission spectrum of the glass does shift somewhat with time and exposure to the air, so that your design should have a safety factor. It is usually a mistake to expose glass filters to severe weathering; in a hostile environment, make the objective (outermost element) out of something more robust (e.g., quartz or borosilicate crown glass).

Glass filters are often called upon to absorb large amounts of optical power, for example, in color-correcting a tungsten bulb with a blue filter to approximate daylight. Glass has poor thermal conductivity, so the temperature does not equilibrate rapidly; this leads to large temperature gradients and consequently large tensile stress in the cooler areas, to the point of shattering the filter. Filter glass therefore is usually tempered to increase the amount of heat that can be dumped into it before it breaks. This is useful but causes severe stress birefringence, which only gets worse with nonuniform heating (see the Schott filter glass catalog). For large heat loads, consider sending the heat away from the filter using a hot or cold mirror, or spreading the load by using a gentle filter to absorb half the heat, then a dense one to absorb the rest.

In long pass filters, the band edge shifts toward the red as the temperature increases, at a rate between 0.02 nm/K for deep UV filters to 0.3 nm/K for NIR filters; it tends to go as

$$\frac{\partial \lambda_c}{\partial T} \approx (5 \times 10^{-7} \text{nm}^{-1}) \lambda_c^2. \tag{5.5}$$

This shift is linear for reasonable temperature excursions, and large enough (hundreds of ppm/°C) to be very obnoxious. There is also a significant shift in passband absorption, which tends to be very large proportionately, since the wings of the exponential are very sensitive to slight changes in kT/e. These shifts are of course sensitive to field angle and NA and so are generally difficult to compensate for in software. If you're trying to do accurate photometry with filters, control their temperature carefully.

Filter glass is usually fluorescent, with a peak about 200 nm to the red of the absorption edge. The usual way of fixing this is to use a series of filters of different cutoff wavelength in series. Unfortunately, the order matters—putting them in the wrong order can cost you a factor of 1000 in leakage. This can be a big effect if you're looking for dim light

in a bright background (e.g., Raman spectroscopy), where it looks just like a light leak, so it'll have you chasing your tail—see Section 10.7.4.

Colored glass filters can have extremely high absorption in their stopbands and are inexpensive; these virtues make up for the gradualness of their absorption versus wavelength compared to interference filters. Unfortunately, due to low sales Schott has reduced the number of filter glasses in their catalog, so that finding a glass just right for your application is significantly more difficult than it once was.

5.5.2 Internal and External Transmittance

Some of the light incident on a dielectric surface is reflected, so that even if the material itself is completely lossless, not all the light hitting a dielectric plate makes it through. We distinguish the two sources of loss by speaking of *internal* and *external transmittance*. Internal transmittance excludes the Fresnel reflections at the surfaces, whereas the external transmittance includes them. For purposes of definition, the filter is assumed thick enough that interference effects and multiple reflections can be ignored.

One benefit of making this distinction is that the dependence of the internal transmittance on the thickness of the element is very simple; it follows *Beer's law*,

$$T_{\text{int}}(\lambda; d) = \exp[-\kappa(\lambda)d], \tag{5.6}$$

which allows us to predict T_{int} of an arbitrary thickness from a single data point:

$$T_{\text{int}}(\lambda; d_2) = [T_{\text{int}}(\lambda; d_1)]^{d_2/d_1}. \tag{5.7}$$

Due to this very strong thickness dependence, the shape of the curve of T_{int} versus d changes with thickness; the width of absorption features increases and the valley transmittance decreases as the element becomes thicker. Filter glass transmittance is usually plotted in terms of the *diabatie*

$$\Theta(\lambda) = 1 - \log_{10}\log_{10}[1/T_{\text{int}}(\lambda; d_0)], \tag{5.8}$$

where $T_{\text{int}}(\lambda; d_0)$ is the internal transmittance of a filter with standard thickness d_0. A plot of diabatie does not change shape with thickness, but merely moves up and down; a common way of presenting the spectral characteristics of filter glass is to plot the diabatie in black on red coordinate axes, then give you a red transparent sheet with the grid and graduations on it. You line up the thickness scale so that the design thickness lines up with the fiducial on the plot, and presto, a plot of internal transmittance versus wavelength for your particular thickness (albeit on a weird vertical scale). Because numerical values of diabatie don't convey much, the scales are labeled with internal transmittance. Neutral density filters are usually thought of in terms of their optical density D,

$$D(\lambda; d) = \log_{10}[T_{\text{ext}}(\lambda; d)]. \tag{5.9}$$

5.5.3 Holographic Filters

Another class of filters is based on holograms. Unlike interference filters, these tend to be available in bandstop types, perhaps 10–20 nm wide, with 40–80 dB (optical) rejection at the stopband center. These devices angle-tune as coatings do, but because of the depth of the null we're talking about here, it leads to a stiffer restriction: you have to use these filters with normally incident collimated beams. A strong beam at the stopband center produces weak surface scatter and stray junk that get through the filter, leading to a doughnut of speckle around the original beam direction. Since they are offset in angle, this is not a critical problem, but you have to put in a baffle after the filter.

5.5.4 Color Correcting Filters

In a tunable system, such as a monochromator driven by a tungsten source, it is often tempting to use color compensation filters, which are fairly gentle colored-glass filters intended to flatten the source spectrum by attenuating the red more than the blue. This should be avoided if possible, for a number of reasons. An optical filter cannot improve the flux of blue light, so that even in the blue, it will decrease the signal-to-noise ratio. The filter response will never be accurately inverse to the instrument function, if for no other reason than that the two change differently with time and temperature, so that a calibration will be necessary anyway. Sometimes there are good technical reasons for using such a filter, for example, a sample or detector that may be damaged by higher intensity in the red, a CCD that blooms badly when saturated, or a digitizer whose dynamic range is insufficient, but these are not as common as might be thought. A slightly more subtle problem is that these filters are mostly designed for color correction with photographic film and do not necessarily make the spectrum even approximately flat. The general idea of *whitening* a signal to improve the SNR is more useful in signal processing—see Sections 13.3.8 and 13.8.10.

5.6 BEAM DUMPS AND BAFFLES

Designing good beam dumps and baffles is a subtle business, which absolutely must be part of the early stages of your instrument design. A common error is to think about baffles last, when even trivial design changes are expensive, and then run for a pricey high performance coating such as Martin Black to fix it.

It must be clearly understood from the outset that the stray light performance of your system is controlled primarily by the geometry rather than by the quality of the black coatings themselves. You can lose a factor of 10^5 by allowing a large expanse of black painted lens barrel, illuminated at grazing incidence, to appear in your detector field of view, and you'll only gain back a factor of 10 or so by replacing the black paint with a fancy and expensive black dendritic finish. The rules are pretty simple.

1. Everything is specular at grazing incidence, so don't allow any grazing bounces to hit your detector.
2. The more illuminated area your detector sees, the more stray light it will receive, so keep the baffles and lens barrel out of the field of view as far as possible.

3. Multiple reflections from dark surfaces will rapidly eliminate stray light, so trap it and then absorb it. Don't allow any one-bounce paths to hit your detector (see rule 2).

4. Sharp edges of baffles will diffract light, so adjust the relative apertures of the baffles to catch it (i.e., later baffles should have slightly smaller inner diameters).

Instruments that search for extrasolar planets need about the best baffles going, which has led to the development of *band-limited* baffles. The idea here is just that of data windowing (see Section 17.4.9), in which a carefully chosen, gradual cutoff of the light leads to greatly reduced diffraction rings and consequently to improved sensitivity at small separations.[†]

5.6.1 What Is a Black Surface?

We call a surface black when it doesn't reflect light. The Fresnel formulas predict significant reflection from any discontinuity in $\tilde{n}$, in either its real or imaginary part. Accordingly, a black surface has an absorption depth of many wavelengths, but much less than its thickness, and is a good index match to the incoming wave. Black surfaces in air are doomed from the outset by the size of the index mismatch at the surface, but if the wave is coming in via a medium such as plastic or glass, the situation is much less dire.

5.6.2 Black Paint

Because of the aforementioned index mismatch, flat black paint in air has a diffuse reflectance of a few percent, which is blacker than TiO_2 but nothing to write home about. (Volume II, Chapter 3 of the *OSA Handbook* has a useful list of black finishes.)

Flat black is useful as a last ditch solution to a scattered light problem, where the stray light is not highly directional. The biggest problem with it is that a major fraction of the light will escape after only one bounce off the black surface, so that the ultimate performance of a flat black paint baffle that is in the detector's field of view is not that much better than that of a single flat black surface. The next biggest is that near grazing incidence, even flat black paint is a quite reasonable reflector. On the other hand, you have to coat the inside of your optical system with something, and flat black is at least less bad than the alternatives.

Internal reflection is a different matter; the improved index match and enforced smoothness of the glass–paint interface improve the qualities of paint enormously (flat black looks shiny from underneath). For example, garden-variety ultraflat black spray paint (Krylon #1602) is a spectacularly good index match to fused quartz, very useful for getting rid of internal reflections from unused areas of quartz prisms. Over the visible, the reflectance of such a quartz–Krylon interface is on the order of 0.01%, which is very impressive for hardware-store spray paint. Remember that paint has environmental limitations and tends to outgas and shed particles.

[†]See, for example, K. Balasubramanian, *Appl. Opt*. **47**(2), 116 (2008).

5.6.3 India Ink

India ink is an aqueous suspension of very fine carbon particles. It is pretty black when dry, but *really* black when liquid, especially if you can get rid of the first-surface reflection—in the visible, the absorption length in India ink is less than 1 cm even at a dilution of $1:10^4$.

5.6.4 Black Anodizing

Anodizing is a surface treatment for aluminum, which consists of electrochemically oxidizing a clean aluminum surface to produce a porous Al_2O_3 (alumina or sapphire) layer. The porosity is inherent—otherwise no current would flow after a very short while, as sapphire is an excellent insulator. The resulting porous matrix is then filled with something else, for example, aluminum hydroxide in the classical anodizing process, or fluoropolymer in some proprietary systems such as Tufram and Polylube. The color comes from dye that is part of the bath. Anodizing is less black than paint because of the high index of sapphire, and in the IR it may not be black at all, since organic dyes do not have the broad absorption spectrum of, say, carbon. Check before relying on its IR performance.

5.6.5 Dendritic Finishes

Coatings made of closely spaced, shiny black peaks or ridges are better absorbers than paint or anodizing. Lockheed Martin makes a dendritic black finish, Martin Black, based on this principle, which is one of a whole range of "designer blacks" (the *OSA Handbook* has a chapter on them). They're useful but far from a complete solution and are very, very expensive. Dendritic finishes tend to reflect a bit more near grazing incidence. Recently, some dendrite-type blacks using oriented carbon nanotubes have got down to about 0.05% reflectance in the visible, but that still won't save a system with lousy baffles.

5.6.6 Black Appliques

There are also a variety of stick-on black finishes, of which the flocked sticky paper sold by Edmund Optics deserves special mention. In the visible, it is comparable in performance to an advanced black coating such as Martin, but costs about 100 times less, and can be cut with scissors. It has a low damage threshold and probably outgasses somewhat due to the adhesive. Because its blackness comes from organic dye, it is much less impressive in the mid-IR, whereas Martin holds up quite well.

5.6.7 Black Plastic

Black plastic is optically similar to glossy black paint, although most of the time its surfaces are not as smooth on small scales. Like black anodizing, some types of black plastic are near-IR transmitting—in the 1970s, some mysterious offset drifts in plastic-packaged op amps were traced to photocurrents induced in the die by light getting through the phenolic plastic. (Modern packages are made of Novolac epoxy, which is very opaque.) If you're in doubt, hold a big sheet of it up to an incandescent light and look through it with an IR viewer (don't use the sun unless you've made sure the plastic is at least opaque enough to be safe for your eyes and your IR viewer).

5.6.8 Black Wax

Carbon black in grease or wax is very black indeed—Wood made his horn with lampblack (candle soot), which was extremely effective as well as convenient. There are a number of hard waxes that can be filled with carbon particles to make a very black material of appropriate refractive index for mounting prisms, especially calcite ones, where the mechanical weakness of wax avoids overstressing the soft crystals. (Apiezon W is a common choice that doesn't outgas.) It isn't a complete solution, though, because black wax doesn't actually stick that well. Prisms are mounted in metal cells that fit the prisms loosely, so the wax is thin and is never loaded in tension. Still, delamination due to mechanical or thermal stress is the major cause of death in calcite prisms.

5.6.9 Black Glass

Various types of very dark colored glass are available, including beer bottles[†] and Schott glasses. These can be UV-epoxied to the faces of prisms or windows in strategic positions, to collect and dissipate stray reflections before they go anywhere. This approach can take a lot more power than spray paint; a thick BK7 flat with a piece of black glass UV-epoxied to the back makes a nice laser reflection attenuator.

Black glass used straight is less satisfactory, as its surface finish is often poor, causing scatter. More subtly, if the laser beam power is high, the temperature gradients in the glass will cause it to become locally convex in the region of high intensity, which will defocus the reflected light. You don't need a kilowatt class laser to see this; 50 mW CW is easily enough. In transmission, this sort of effect is called thermal lensing.

5.6.10 Designing Beam Dumps and Light Traps

Assuming that you've designed the system sensibly, stray light will have to make at least one bounce to get to the detector. Thus controlling stray light involves two steps: reducing the amount of illuminated area in the field of view of the detector, and reducing the illumination intensity there. All that stray light has to go somewhere, and that somewhere is usually a baffle of some sort. We send unwanted beams into a *beam dump* and corral ambient light and scatter into a *light trap*. The two look more or less the same.

The best way to design beam dumps is to use *shiny* black surfaces, where the surface reflection is specular. A specular reflection can be directed onto another black surface, and another and another. . . . With care, you can make the light take many bounces before it can escape. The job is foremost to *trap* the light, and then to dispose of it.

5.6.11 Wood's Horn

The original beam dump is Wood's horn,[‡] a gently curved and tapered tube of glass coated outside with lampblack or black paint, shown in Figure 5.13a. It is a really good design, which works well over a reasonable range of incidence angles. The gentle taper traps the specular reflections, and due to the curved shape, most of the diffusely scattered light also has to make multiple bounces before escaping. Because the beam rattles around

[†]Beer's law is named after Dr. Beer.
[‡]Named after Robert Williams Wood, debunker of *N*-rays (*Nature* **70**, 530–531 (1904)), pioneer of grating spectroscopy, and author of *How to Tell the Birds from the Flowers*, among other stellar contributions.
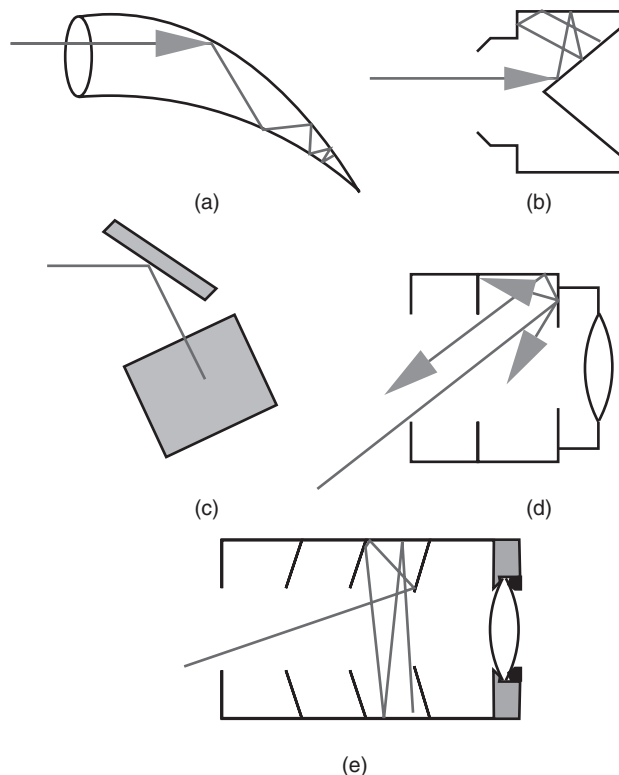
Figure 5.13. Assorted beam dump and baffle designs: (a) Wood's horn, (b) cone dump, (c) black glass at Brewster's angle, (d) knife-edge baffles, and (e) barbed baffles. Designs (a)–(c) use shiny black surfaces, and (d) and (e) shiny or flat black.

between surfaces making some angle with each other, it tends to be sent back out after some number of reflections, so the length of the horn has to be at least a few times its diameter.

5.6.12 Cone Dumps

A more convenient beam dump is the conical type, which fits optical breadboards and erector sets such as Microbench. As shown in Figure 5.13b, the cone absorbs most of the light and directs the rest into a simple trap arrangement that tends to confine the light. Light has to make at least three bounces from shiny black surfaces to escape, and most makes many more. These are easy to build in the lab if you have a lathe.

5.6.13 Black Glass at Brewster's Angle

You can combine black glass with Brewster angle incidence to get rid of an unwanted collimated beam (e.g., the residual pump laser beam in a photoacoustic measurement). As shown in Figure 5.13c, the first piece of black glass at Brewster's angle gets rid of one polarization, and the second one, at Brewster's angle for the other polarization (partly reflected from the first one), completes the job. This approach can take more peak power

than Wood's horn, but is restricted to well-collimated beams coming from one particular direction, and requires attention to the control of surface scatter from the black glass.

5.6.14 Shiny Baffles

A barbed pattern, with the barbs pointing toward the light slightly, is probably the best sort of baffle for long lens barrels and other narrow cylinders. Make the channels narrow enough that a beam needs several reflections to escape. Coat the inside with shiny black paint or make it from a shiny black material. The disadvantage of barbed patterns is that the light will eventually be reflected back out, and that the number of reflections required for this is a strong function of the angle of incidence.

5.6.15 Flat Black Baffles

In some situations, strong stray light can come from far off axis and may enter at any angle, as with sunlight in outdoor applications, so this angular dependence is inconvenient. In cases like this, we may resort to flat black surfaces and just reduce the illuminated area in the field of view. Optical design packages have stray light analysis that relies on the bidirectional reflectance distribution function (BRDF), which predicts the amount of light scattered into $\mathbf{k}_2$ from $\mathbf{k}_1$. Do yourself a favor and design decent baffles.

An example of a flat black surface approach that works well is the knife-edge baffle, consisting of a series of black apertures lying normal to the axis of the optical system. Knife edges are easy to fabricate, being planar structures. The inside diameters decrease slightly coming toward the detector, so that ideally the earlier baffles in the series are out of the detector's field of view entirely, and thus light escaping from the flat black surfaces must either go back out the objective or hit another baffle. You do need to make the edges sharp, though, because if they're too blunt you'll have nice grazing incidence reflectors sending light into your detector. Knife edge baffles are lightweight and highly effective when properly designed.[†] Figure 5.13 shows some popular baffles and beam dumps.

5.6.16 Combinations

If the optical system is sufficiently long and narrow that the first few knife-edge baffles are really out of the detector FOV, you can use shiny black surfaces there and apply simple graphical ray tracing (starting from the detector) to estimate what the stray light intensity will be. A ray hitting baffle n will be near enough to normal incidence that it will rattle around between baffles n and $n - 1$ several times before exiting, which will get rid of it very effectively. If baffle n is smaller in aperture, and none of baffle $n - 1$ is visible to the detector, then a lot of the escaping light will exit the way it came in, which helps. Subsequent baffles, which may be in the FOV, can be flat black if necessary.

[†]See, for example, A. Buffington, B. V. Jackson, and C. M. Korendyke, Wide-angle stray-light reduction for a spaceborne optical hemispherical imager. *Appl. Opt*. **35**(34), 6669–6673 (1996).

5.7 WHITE SURFACES AND DIFFUSERS

A white surface is one that scatters incident light of all colors efficiently. We usually want diffuse white surfaces, whose scatter pattern is nearly Lambertian. White surfaces are made by suspending small particles of high index, nearly lossless dielectric in a low index medium, to scatter light as strongly and as often as possible.

5.7.1 Why Is It White?

When light enters a white surface, it is scattered in all directions; thus much of it does a random walk in the material. When it encounters the surface, most of it will escape (all of it, apart from Fresnel losses and TIR). Mathematicians call this phenomenon *gambler's ruin*—due to that boundary, eventually all your money diffuses out of your pocket. Considering how many times it will be scattered, and the geometric dependence of intensity on the number of scatterings, any absorption will be enormously enhanced; the same metals we use for mirrors appear black in powder form. Some white surfaces are better at depolarizing light than others, so measure yours carefully if it matters to you.

5.7.2 Packed Powder Coatings

The best diffuse, high reflection surface is a packed powder of pure TiO_2, $Ba(SO_4)$, or MgO in air. Their reflectance is over 99% through most of the visible, and it is highly Lambertian (TiO_2's drops off badly near 410 nm, and in commercial grades it also tends to fluoresce). What's more, unlike paint it is easily renewed if it gets dirty. It is primarily useful as a reflectance standard, because the surface is easily disturbed and really has to lie nearly horizontal. Barium sulfate's claim to fame is that its reflectivity is very constant with time and very flat with wavelength. Packed polytetrafluoroethylene (PTFE) powder will stick to surfaces well enough to use it in integrating spheres and has similar reflectance properties. Because of its lower refractive index, it needs a thicker section (at least 6 mm or so) to reach its peak reflectance, but can achieve values above 0.996 in the visible, and maintains its properties into the UV and NIR.[†]

5.7.3 Barium Sulfate Paint

The most popular diffuse white coating is barium sulfate paint (available from Edmund Optics as "Munsell white reflectance coating"). Paint is a collection of various particles in a dielectric binder. There is a front-surface reflection from the binder, which makes painted surfaces non-Lambertian, though they're closer than most other things. Compared with $BaSO_4$ powder in air, the dielectric is lossier, and the index mismatch at the surfaces is smaller, so the total reflectivity is also lower—about 98% in most of the visible and NIR. Regular white paint is loaded with TiO_2 and is closer to 90% in the visible. Barium sulfate paint is especially good in places like the interior of integrating spheres, where you need a nearly Lambertian reflector that's very opaque in a fairly thin layer (1–2 mm), and there's a lot of nonhorizontal area to cover.

[†]Victor R. Weidner and Jack J. Hsia, Reflection properties of pressed polytetrafluoroethylene powder. *J. Opt. Sci. Am*. **71**, 7 (July 1981).

5.7.4 Spectralon

Spectralon is a sintered PTFE sold by Labsphere, with properties similar to packed PTFE powder. It can be machined into odd shapes and is stable and cleanable. (Avian Technology sells similar stuff as Fluorilon-99W.) It is highly reflective and, although not so Lambertian as fine powders, it is very convenient for applications needing high efficiency diffuse reflectors that can absorb some punishment. The stuff is very expensive, though, so don't go milling an integrating sphere from a solid block of it.

5.7.5 Opal Glass

Opal glass is very inefficient at transmission (1%) but very Lambertian. It is used only for those applications for which diffuse illumination is vital. It produces very small speckles when used with lasers.

5.7.6 Magic Invisible Tape

Matte finish translucent tape is intended for mending torn papers, but it works pretty well as a diffusing material for light duty use, for example, putting a piece on the entrance slit of a spectrometer to fix spectral artifacts due to a weird pupil function, or to homogenize the ugly beam patterns of LEDs and liquid light guides. It won't take much power, and it leaves a slight residue behind, but it lasts for years, so it's just the right medicine sometimes.

5.7.7 Integrating Spheres

Light reflected from a white coating loses most of its directional information in a single bounce. A closed cavity with sufficiently high reflectance can bounce the light dozens of times before absorbing it, so that the illumination of a point on the wall becomes Lambertian to high accuracy; this is the idea of an integrating sphere. There are two main applications: measurement of optical power and building Lambertian light sources. The photon efficiency of an integrating sphere is a great deal higher than that of opal glass, so you can make a good light source by putting a small bulb inside the sphere, screened by a small white shield so that no unscattered light can reach the exit hole. The hole has to be fairly small, no more than 1/6 of the sphere diameter, to get the best performance. The same homogenizing property makes integrating spheres the best optical power sensors available; a photodiode in place of the bulb (still behind the shield) will see almost the same total flux regardless of the incident angle of the optical beam, assuming it isn't vignetted by the aperture, and furthermore the spatial and angular variations of the responsivity is homogenized out. Residual angular variation is at the 0.1% level unless the ports are too large.

The average number of bounces required for light to escape is equal to the total area of the apertures (projected on the sphere) divided by the area of the sphere. Assuming this is small, we can treat it like reflection loss, so by the geometric series formula,

$$\eta = \frac{P_{\text{out}}}{P_{\text{in}}} \approx \frac{A_{\text{out}}}{4\pi r^2(1 - R) + (A_{\text{out}} + A_{\text{in}})R}, \tag{5.10}$$

where R is the reflectance of the coating, r is the inside radius of the sphere, and A_{out} and A_{in} are the areas of the output and input ports (actually the areas of their projections on

the sphere). This doesn't take account of the few bounces it takes for the light to become Lambertian inside the sphere, the effect of baffles inside the sphere, or the deviation of the steady state illumination from a Lambertian condition due to the losses through the ports, all of which are generally small effects. For a perfectly reflecting sphere with equal sized ports, $\eta = 0.5$, and in real spheres, it is generally no more than 0.3 and is commonly much lower.

Similarly, a δ-function light impulse will be spread out into a roughly exponential pulse of time constant

$$\tau \approx \frac{4r}{3c} \frac{1}{(1 - R) + (A_{\text{out}} + A_{\text{in}})R/(4\pi r^2)}, \tag{5.11}$$

which is on the order of 10–50 ns for most spheres.

From (5.10), we see that the efficiency of a sphere is a very strong function of the coating reflectance, particularly if the coating is very good. Spheres made with the very highest reflecting coatings are therefore somewhat less stable with time and conditions than those made with ordinary white paint; on the other hand, they do a better job of diffusing the light and waste fewer photons. Keep your integrating spheres clean, plug any unused apertures with the white caps provided, and use them in their intended wavelength interval. This sensitivity can also be used to advantage in multipass measurements: see Section 10.6.5 for an example. The many bounces taken by a typical photon before it is absorbed or lost unfold to quite a long optical path, as we saw, and this can be very helpful in reconciling the demands of fast pulses to the capabilities of photodiodes, as in Section 3.5.4.

5.7.8 Ping-Pong Balls

You can make a rough-and-ready integrating sphere from a ping-pong ball. Paint the outside white, drill two holes at $120°$ from each other, and put a photodiode in one of the holes. This is good enough to show some of the advantages of spheres, but not for real measurements. (Ping-pong balls are also good scatterometers—see Section 9.8.7.)

5.7.9 Ground Glass

Ground glass is much more efficient at light transmission than opal glass but has a big specular component as well as a diffuse component. Use it where the specular component is not a critical problem, because besides being cheaper than opal glass, it is dramatically more efficient (30–70% vs. 1%). Because of TIR, it matters which way round you put the ground side; light crossing the ground surface from inside the glass is somewhat more diffuse but dimmer. Objects whose light scattering occurs at a surface tend to produce a constant spread of u and v, centered on the unscattered $\mathbf{k}$ vector. Thus the angular spectrum is not independent of the incidence angle, which complicates diffuser design. If you have something like a video projector with a short focal length lens, shining on ground glass, it will scatter light only a bit around the original $\mathbf{k}$ vector, so at the edges most of the light will still be going up and away rather than straight out, as you would probably want. Software can correct for this at a single viewing angle, but not for everyone in the room. This sounds like a job for a field lens (see Section 12.3.14)—the big Fresnel lens of Figure 5.14. straightens out the light before it hits the ground glass, which makes the brightness look much more uniform with viewing angle (though no closer to Lambertian than before).
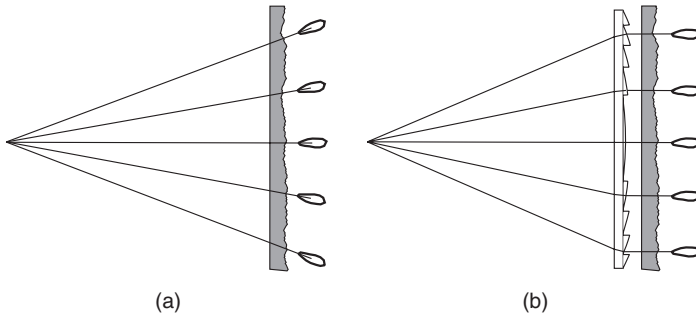
Figure 5.14. Ground glass and other mild diffusers tend to scatter light into a cone about the incident ray direction as in (a). Adding a Fresnel field lens as in (b) can improve the apparent uniformity.

5.7.10 Holographic Diffusers

A better controlled version of ground glass can be made with a holographic element, the *holographic diffuser*. These are touted as being useful for laser beam shaping, but in reality the strong speckle they produce limits their application to low coherence sources such as LEDs. One very good application is to homogenize the output of fiber bundle illuminators. Nowadays you can get holographic diffusers that are nearly as Lambertian as opal glass, or have other angular patterns such as top-hat or square, without the high losses of opal glass diffusers. Not all holographic diffusers have the other special properties of opal glass (e.g., independence of illumination angle).

5.7.11 Diffusers and Speckle

A light source such as a HeNe laser, which is highly coherent in both space and time, is very difficult to use with diffusers. Shining a HeNe into an integrating sphere produces an optical field that is very Lambertian on a broad-area average, but that has very strong small-scale structure called *speckle*. All rough surfaces illuminated with lasers produce speckles that are a complicated function of position, but whose size is characteristic of the material and of the size of the illuminated region. More diffuse materials produce smaller speckles; the angular extent of the smallest ones is on the order of λ/d, where d is the incoming beam diameter. Speckle consists of a mass of unorganized interference fringes, caused by the coherent summation of fields from everywhere in the sphere. At each point, these random reflections produce a certain optical amplitude and phase in each polarization component, which vary all over the place. The best diffusers, such as integrating spheres, produce speckles with characteristic size $\lambda/2$. Due to speckle statistics, the relative standard deviation of the photocurrent will be on the order of $\sqrt{\text{(PD area)(speckle area)}}$, which isn't that small, and any vibration will smear that out into a huge noise PSD in the low baseband. Thus diffusers aren't always the way to get good measurements, at least with lasers. See Section 2.5.1 for more discussion.