

Polarization

He flung himself on his horse and rode madly off in all directions.

—Stephen Leacock, *Gertrude the Governess*

6.1 INTRODUCTION

Optical polarization is the main way the vector wave nature of light manifests itself in practical problems. We've encountered plane waves, which are always perfectly polarized, and the Fresnel formulas, which predict the intensity and polarization of plane waves leaving a dielectric surface. Here we go into the measurement and manipulation of polarization, and how not to get in trouble with it. Polarization components such as retarders and Faraday rotators are mysterious to lots of people, but are actually fairly simple devices unless you try to get deeply into the physics of how they do what they do. Being practical folk, we'll stick with their phenomenology and keep their inner workings pretty well out of it.

The major uses of polarization components in optical systems are to control reflections, as in sunglasses and fiber isolators, and to split and combine beams without the heavy losses caused by ordinary beamsplitters.

6.2 POLARIZATION OF LIGHT

6.2.1 Unpolarized Light

If you send thermal light through an analyzer,[†] twist the control ring as you may, the same proportion of the light comes through. This remains true if you put any sort of lossless polarization device ahead of it; a wave plate or a Faraday rotator doesn't change the polarization at all. Thus we say that thermal light is *unpolarized*. This is a poser, because we know that any optical field can be decomposed into plane electromagnetic

[†]Analyzers and polarizers are physically identical, but an analyzer is thought of as detecting the polarization state produced by the polarizer—in communications terms, the analyzer is part of the receiving section, and the polarizer is part of the transmitting section.

waves. Since all such waves are perfectly polarized, how can thermal light be unpolarized?

The key is that we're really measuring the *time-averaged* polarization rather than the instantaneous polarization. The light at any point at any instant does in fact have a well-defined \mathbf{E} vector, because if it didn't, its energy density would be 0. In an unpolarized field, though, the direction of \mathbf{E} varies extremely rapidly with time, changing completely in a few femtoseconds in the case of sunlight. Thinking in terms of modulation frequency (see Section 13.3), the polarization information is not concentrated at baseband the way it is with lasers, but instead is smeared out over hundreds of terahertz of bandwidth. It is spread so thin that even its low frequency fluctuations are hard to measure.

In \mathbf{k} -space terms, the polarizations of different plane wave components are completely uncorrelated, for arbitrarily close spacings in \mathbf{K} . This is in accord with the entropy-maximizing tendency of thermal equilibrium—any correlation you could in principle use to make energy flow from cold to hot is always 0 in thermal equilibrium.

6.2.2 Highly Polarized Light

If we pass thermal light through a good quality polarizer, we get highly polarized thermal light. The plane wave components are still uncorrelated in phase but are now all in the same polarization state. If such light does not encounter any dispersive birefringent elements, its polarization state may be modified but it will remain highly polarized. Its polarization can be changed achromatically with TIR elements such as Fresnel rhombs, so that we can have thermal light with a well-defined circular or elliptical polarization.

6.2.3 Circular Polarization

We've encountered circular polarization before, but there's one property that needs emphasizing here, since so many useful polarization effects depend on it: the helicity changes sign on reflection. Left-circular polarization becomes right circular on reflection, and vice versa— \mathbf{E} keeps going round the same way, but the propagation direction has reversed, so the helicity has reversed too. This is also true of ordinary screw threads viewed in a mirror, so it's nothing too mysterious. Although linear polarization can be modified on oblique reflection from a mirror (if \mathbf{E} has a component along the surface normal), circular polarization just switches helicity, over a very wide range of incidence angles.[†] Since linear polarization can be expressed in terms of circular, this should strike you as odd—there's a subtlety here, called topological phase, that makes it all come out right in the end.

6.2.4 An Often-Ignored Effect: Pancharatnam's Topological Phase

When light traverses a nonplanar path, for example, in two-axis scanning, articulated periscopes, or just piping beams around your optical system, its polarization will shift.

[†]If the reflection occurs at a dielectric interface (where $r_p \neq r_s$), the polarization will become elliptical, at θ_B the ellipse degenerates into linear polarization, and beyond θ_B , the helicity no longer reverses. (Why?)

For reflection off mirrors, this isn't too hard to see: since \mathbf{E} is perpendicular to \mathbf{k} , a mirror whose surface normal has a component along \mathbf{E} will change \mathbf{E} . Make sure that you follow your polarization along through your optical system, or you may wind up with a nasty surprise.

A much more subtle fact is that the same is true for any system where light travels in a nonplanar path (e.g., a fiber helix). Left- and right-circular polarizations have different phase shifts through such a path, giving rise to exactly the same polarization shift we get from following the mirrors; this effect is known as Pancharatnam's topological phase[†] and is what accounts for the puzzling difference in the polarization behavior of linear and circularly polarized light upon reflection that we alluded to earlier (the corresponding effect in quantum mechanics is Berry's phase, discovered nearly 30 years after Pancharatnam's almost-unnoticed work in electromagnetics). This sounds like some weird quantum field effect, but you can measure it by using left- and right-hand circular polarized light going opposite ways in a fiber interferometer.[‡] These polarization shifts are especially important in moving-mirror scanning systems, where the resulting large polarization shift may be obnoxious.

It sounds very mysterious and everything, but really it's just a consequence of spherical trigonometry; the \mathbf{k} vector is normal to a sphere, and \mathbf{E} is tangent to the sphere throughout the motion; depending on how you rotate \mathbf{k} around on the surface, \mathbf{E} may wind up pointing anywhere. Equivalently, 2×2 rotation matrices commute, but 3×3 ones don't.

If you follow your \mathbf{k} vector around a closed loop enclosing a solid angle Ω , the relative phase of the right- and left-circular polarizations gets shifted by

$$\Delta\phi = \pm 2\Omega. \quad (6.1)$$

6.2.5 Orthogonal Polarizations

We often describe two polarization states as *orthogonal*. For linear polarizations, it just means perpendicular, but what about circular or elliptical ones? The idea of orthogonal polarizations is that their interference term is 0, that is,

$$\mathbf{E}_1 \cdot \mathbf{E}_2^* = 0. \quad (6.2)$$

Two elliptical polarizations are thus orthogonal when their helicities are opposite, their eccentricities equal, and their major axes perpendicular (i.e., opposite sense of rotation, same shape, axes crossed). It's an important point, because as we'll see when we get to the Jones calculus in Section 6.10.2, lossless polarization devices do not mix together orthogonal states—the states will change along the way but will remain orthogonal throughout. One example is a quarter-wave plate, which turns orthogonal circular polarizations into orthogonal linear polarizations, but it remains true even for much less well-behaved systems such as single-mode optical fibers.

[†]S. Pancharatnam, Generalized theory of interference and its applications. Part 1. Coherent pencils. *Proc. Indian Acad. Sci.* **44**, 2247–2262 (1956).

[‡]Erna M. Frins and Wolfgang Dultz, Direct observation of Berry's topological phase by using an optical fiber ring interferometer. *Opt. Commun.* **136**, 354–356 (1997).

6.3 INTERACTION OF POLARIZATION WITH MATERIALS

6.3.1 Polarizers

A polarizer allows light of one polarization to pass through it more or less unattenuated, while absorbing or separating out the orthogonal polarization. Any effect that tends to separate light of different polarization can be used: anisotropic conductivity, Fresnel reflection, double refraction, walkoff, and the different critical angles for *o*- and *e*-rays (related to double refraction, of course).

Polarizers are never perfectly selective, nor are they lossless; their two basic figures of merit at a given wavelength are the loss in the allowed polarization and the *open/shut ratio* of two identical polarizers (aligned versus crossed) measured with an unpolarized source, which gives the polarization purity. The best ones achieve losses of 5% or less and open/shut ratios of 10^5 or even more.

6.3.2 Birefringence

The dielectric constant $\underline{\epsilon}$ connects the electric field \mathbf{E} with the electric displacement \mathbf{D} ,

$$\mathbf{D} = \underline{\epsilon}\mathbf{E}. \quad (6.3)$$

For a linear material, $\underline{\epsilon}$ is a tensor quantity (in isotropic materials the tensor is trivial, just ϵ times the identity matrix).[†] (See also Section 4.6.1.) Tensors can be reduced to diagonal form by choosing the right coordinate axes; the axes that diagonalize $\underline{\epsilon}$ are called the *principal axes* of the material; symmetry requires that they be orthogonal in this case. (The refractive index also of course may depend on polarization but is not a tensor, because it does not express a linear relationship.)

Some common birefringent optical materials are crystalline quartz, sapphire, calcite (CaCO_3), and stretched plastic films such as polyvinyl alcohol (PVA) or polyvinylidene chloride (Saran Wrap). All these, along with most other common birefringent materials, are *uniaxial*[‡]; two of their three indices are the same, $\epsilon_x = \epsilon_y = \epsilon_{\perp}$; light polarized in the plane they define is an ordinary ray (*o*-ray), so called because it doesn't do anything strange. The third index, which defines the *optic axis*, may be larger (positive uniaxial) or smaller (negative uniaxial) than the *o*-ray index; if \mathbf{E} has a component along the optic axis direction, strange things occur, so that the beam is called an *e*-ray, for "extraordinary." Things get stranger and less relevant for absorbing birefringent materials and for biaxial ones, so we'll stick with the important case: lossless uniaxial materials.

Electromagnetic fields are *transverse*, which in a uniform medium means that for a plane wave, \mathbf{E} , \mathbf{H} , and \mathbf{k} are always mutually perpendicular, and that the Poynting vector \mathbf{S} always lies along \mathbf{k} . (The Poynting vector generally defines the direction the energy

[†]Landau and Lifshitz, *The Electrodynamics of Continuous Media*, has a lucid treatment of wave propagation in anisotropic media, which the following discussion draws from.

[‡]Less symmetric materials may be biaxial, that is, have three different indices, and in really messy crystal structures, these axes need not be constant with wavelength. Biaxial crystals exhibit some weird effects, such as *conical refraction* (see Born and Wolf).

goes in; that is, it's the propagation axis of the beam as measured with a white card and ruler).[†]

Neither of these things is true in birefringent materials, where we have only the weaker conditions that \mathbf{D} , \mathbf{B} , and \mathbf{k} are mutually perpendicular, as are \mathbf{E} , \mathbf{H} , and \mathbf{S} . For instance, the actual index seen by the e -ray changes with angle, unless the light propagates in the plane defined by the ordinary axes, for only then can \mathbf{E} lie exactly along the optic axis. The propagation vector \mathbf{k} defines an ellipsoid (where x , y , and z are the principal axes),

$$\frac{k_x^2}{\epsilon_x} + \frac{k_y^2}{\epsilon_y} + \frac{k_z^2}{\epsilon_z} = k_0^2. \quad (6.4)$$

The refractive index $n = k/k_0$ experienced by a given e -ray varies with its propagation direction. The extreme values of n_e are n_\perp (the o -ray index n_o) when \mathbf{k} is along the optic axis and n_\parallel when \mathbf{k} is normal to the optic axis. There is a lot of sloppiness in the literature, with n_\parallel often being referred to as n_e , whereas n_e really varies between n_\perp and n_\parallel . Light encountering the surface of a birefringent material is split apart into two linearly polarized components going in different directions. Phase matching dictates that \mathbf{k}_\perp is preserved across the boundary. The o -ray behaves just as if the material were isotropic with an index of n_\perp , so that's easy— \mathbf{S} is parallel to \mathbf{k} .

Determining \mathbf{k} and \mathbf{S} for the extraordinary ray is straightforward. The direction and magnitude of \mathbf{k}_e can be found from (6.4) and the phase matching condition. Once \mathbf{k}_e is known, the direction of \mathbf{S} can be found easily; it lies in the plane defined by \mathbf{k} and the optic axis, and the angles θ_k and θ_S separating the optic axis from \mathbf{k}_e and \mathbf{S} obey

$$\tan \theta_S = \frac{\epsilon_\perp}{\epsilon_\parallel} \tan \theta_k. \quad (6.5)$$

Remember, though, that the phase term is still $\exp(i\mathbf{k} \cdot \mathbf{x})$ —stick with this and don't get confused by trying to calculate propagation distance along \mathbf{S} and multiplying by $n_e k_0$ or something like that.

If light travels directly down the optic axis, \mathbf{E} has no component along it, so the material appears to be isotropic. This is useful where the birefringence is obnoxious, for example, “ c -axis normal” sapphire windows used for their strength and chemical inertness.

6.3.3 Retardation

Since the phase velocity of light in a material is c/n , the e - and o -rays propagate at different phase velocities, so the two linear polarization components with \mathbf{k} along z will be phase shifted with respect to each other by an amount δ , where

$$\delta = (n_e - n_o)k_0 z. \quad (6.6)$$

[†]Care is needed in identifying $\mathbf{E} \times \mathbf{H}$ with the local energy flux: Poynting's theorem applies to the integral of $\mathbf{S} \cdot d\mathbf{A}$ over a closed surface, or equivalently with the volume integral of $\nabla \cdot \mathbf{S}$. That means that \mathbf{S} is nonunique in much the same way as the magnetic vector potential—Poynting's theorem still holds if we add to \mathbf{S} the curl of any arbitrary vector field. It usually works.

Unless the incoming beam is a pure *e*- or *o*-ray, this will change the resulting polarization (as we saw in Section 1.2.8). This phenomenon is called *retardation* and is the basis for wave plates. Retardation is usually specified in nanometers, since it is a time delay Δt that causes a phase shift $\delta = \omega \Delta t$, in contrast to a reflection phase as in a Fresnel rhomb, which is almost wavelength independent. (In other words, retarders are wavelength dependent even when the material has no dispersion.)

6.3.4 Double Refraction

An oblique beam entering such a material from an isotropic medium splits into two beams, because the different refractive indices give different angles of refraction by Snell's law. This phenomenon is called *double refraction* (which is what birefringence means, of course).

6.3.5 Walkoff

Besides double refraction, birefringent materials exhibit *walkoff*, as shown in Figure 6.1. Although the \mathbf{k} vector behaves normally in a birefringent material, the Poynting vector does not; the energy propagation of the *e*-ray is not parallel to \mathbf{k} , but lies in the plane defined by \mathbf{k} and the optic axis, somewhere between them, so that the *e*-ray seems to walk off sideways. This weird effect arises because the Poynting vector is parallel to $\mathbf{E} \times \mathbf{H}$. The tensor character of ϵ prevents \mathbf{E} from being perpendicular to \mathbf{k} ,[†] and the




Figure 6.1. Polarizers based on beam walkoff: (a) simple walkoff plate or beam displacer and (b) the Savart plate, a symmetrical walkoff plate.

[†]Unless \mathbf{D} is an eigenvector of ϵ , that is, is a pure *o*-ray or lies along the optic axis.

cross-product relation then forces \mathbf{S} to be not along \mathbf{k} . This effect has nothing to do with double refraction; instead, it's a spatial analogue of the phase velocity/group velocity distinction for a pulse of light. A general beam normally incident on a planar slab of birefringent material will split apart into two beams going in different directions. Double refraction can't cause this directly, since at normal incidence no refraction occurs. Oblique beams walkoff, as well, but is less obvious then. Now you know why it's called the extraordinary ray.

Aside: Defuzzing Filters. Very thin walkoff plates, usually LiNbO_3 , are often used in CCD cameras to reduce the disturbing moiré patterns due to the way the pixels are arranged in color cameras (see Section 3.9.14). Two walkoff plates mounted at 90° to one another, with a $\lambda/4$ plate in between, split an image point into an array of four points, thus multiplying the OTF of the camera by $\cos(u dx) \cos(v dy)$, where dx and dy are the shift distances. This apodization rolls off the OTF to zero at frequencies where the moiré patterns are objectionable. (Sometimes quartz is used, but it has to be thicker, which causes more aberration.)

6.3.6 Optical Activity

A birefringent material has different refractive indices for different linear polarizations. A material that has some inherent helicity, such as a quartz crystal or a sugar solution, has different indices for different circular polarizations (helical antennas respond more strongly to one circular polarization than to the other). The different coupling between the bound electrons and the radiation field gives rise to a slightly different index of refraction for the two helicities. As a linearly polarized wave propagates in such a medium, \mathbf{E} changes direction; if there is no birefringence, \mathbf{E} describes a very slow helix as it propagates. This is called *optical activity* or *circular birefringence*.

Noncentrosymmetric crystals such as α quartz and tellurium dioxide may exhibit both optical activity and birefringence; this combination makes the polarization effects of a random hunk of crystal hard to predict.

If you put a mirror on one side of a piece of isotropic but optically active material (e.g., a cuvette of sugar water), the linear polarization that comes out is exactly the same as the one that went in; the rotation undoes itself. This is because the helicity reverses itself on reflection—each component crosses the material once as left circular and once as right circular, so that their total delays are identical, and the original polarization direction is restored.

The effects of optical activity are fairly weak but highly dispersive; for a 90° rotation in α quartz, you need 2.7 mm at 486 nm and 12 mm at 760 nm; this is around 100 times weaker than the effect of birefringence, so it dominates only when the light is propagating right down the optic axis. The dispersion is occasionally useful (e.g., in separating laser lines with low loss), but since it's an order of magnitude higher than a zero-order wave plate's, optical activity isn't much use with wideband light except for making pretty colors.

Due to the columnar morphology of coatings (see Section 5.4.7), it is possible to make artificial circular birefringent coatings by evaporation. The substrate is held at an angle to the source and slowly rotated about the source–substrate axis, producing helical columns that are highly optically active.

6.3.7 Faraday Effect

Another effect that leads to the slow rotation of \mathbf{E} is the Faraday or magneto-optic effect, which is often confused with optical activity because the effects are superficially very similar. Terbium-doped glasses and crystals such as terbium gallium garnet (TGG), immersed in a magnetic field, rotate the polarization of light propagating parallel to \mathbf{B} by

$$\Theta = VBl, \quad (6.7)$$

where Θ is the rotation angle, l is the path length, B is the axial magnetic field, and V is a material property called the Verdet constant. The difference here is that there is a special direction, defined by \mathbf{B} . Heuristically, if you imagine that the application of \mathbf{B} starts a bunch of bound currents going around in circles, then what matters is not the helicity but whether \mathbf{E} is rotating the same way as the currents or not, because the dielectric susceptibility will be different in the two cases.

The key point is that the rotation direction does not change on reflection. If we put our mirror at one side of the magneto-optic material, \mathbf{E} keeps going round the same way on both passes, so the helicity change on reflection does not make the delays equal; Faraday rotation doubles in two passes, instead of canceling—it is said to be *nonreciprocal*. This property allows us to build *optical isolators*, which allow light to propagate one way but not the other, and Faraday rotator mirrors, which help undo the polarization nastiness of optical fibers.

6.4 ABSORPTION POLARIZERS

Some materials exhibit polarization-selective absorption, as in Polaroid sunglasses. They do it by anisotropic conductivity, which is what you'd expect given the close relationship between conductivity and the imaginary part of n .

6.4.1 Film Polarizers

Film polarizers are made of anisotropically conductive polymer: stretched polyvinyl alcohol (PVA) doped with iodine. They work throughout the visible, but deteriorate in the infrared, and are useless in the ultraviolet since PVA is strongly absorbing there. There are several different kinds, for different wavelength intervals, but the good ones absorb about 20–40% of the allowed polarization and have open/shut ratios of 10^4 . The selectivity of older types used to degrade significantly in the blue, but the newer ones are much better. Their wavefront quality is relatively poor (about like window glass, $2\lambda/\text{inch}$), and they have a very low damage threshold, only 1 W/cm^2 or so.

In order not to be limited by the wavefront wiggles, put the polarizer near the image. An image formed in thermal light has very small phase correlations between points to begin with (since the phase of object points further than λ/NA apart is essentially uncorrelated), so phase wiggles are not much of a worry at an image.

6.4.2 Wire Grid Polarizers

Wire grids, which despite their name are arrays of very thin, closely spaced, parallel wires, function well in the mid- and far-infrared, but the difficulty of making the pitch

fine enough prevents their use in the visible—such a structure is obviously a diffraction grating, so the pitch has to be fine enough that the first diffracted order is evanescent (see Section 7.2). Their open/shut ratios are usually about 10^2 , and they absorb or reflect about 50% of the allowed polarization. They also reflect some of the rejected polarization, but how much depends on the metal and the geometry. Shorter-wavelength grids are usually lithographically deposited, so you have to worry about substrate absorption as well.

6.4.3 Polarizing Glass

A development of the wire grid idea is the dichroic[†] glass Polarcor, made by Corning. It is an optical glass with small metallic silver inclusions. During manufacture, the glass is stretched along one axis, which transforms the inclusions into small whiskers, aligned with the axis of the stretch. These whiskers function like little dipole antennas and are highly absorbing in a relatively narrow band. At present, Polarcor is best used in the near-infrared (out to $1.6\ \mu\text{m}$) but is available down to 600 nm in the visible (transmittance deteriorates somewhat toward short wavelengths). It has excellent transmission in one polarization (70–99%), and excellent extinction in the other ($\approx 10^{-5}$), so that its open/shut ratio is comparable to that of crystal polarizers. It has a wide ($\pm 30^\circ$) acceptance angle and good optical quality—though there are a few more striae than in ordinary optical glass, as one would expect. Polarcor's laser damage threshold is lower than calcite's—if you hit it too hard, the silver grains melt and lose their shape. The threshold is around $25\ \text{W}/\text{cm}^2$ CW, or $0.1\ \text{J}/\text{cm}^2$ pulsed.

6.5 BREWSTER POLARIZERS

At intermediate angles of incidence, reflections from dielectric surfaces are fairly strongly polarized. At Brewster's angle, R_p goes to 0, and for glass $R_s \approx 10\%$ per surface. The effect in transmission is not strong enough to qualify as a polarizer unless it's used intra-cavity, but can be enhanced by doing it many times. A pile of microscope slides at 55° or so makes a moderately effective polarizer, and (as we saw in Section 5.4.4) a $(HL)^m H$ stack of dielectric films can be highly effective.

6.5.1 Pile-of-Plates Polarizers

Assuming that the light is low enough in coherence that etalon fringes can be ignored, m glass plates stacked together and oriented at θ_B will attenuate the s -polarized light by a factor of 0.8^m , which for 31 plates amounts to 10^{-3} , ideally with no loss at all in the p polarization. This nice property is of course degraded as θ_i departs from θ_B , but it's useful over a reasonable angular range. The transmitted wavefront fidelity of a pile-of-plates polarizer is poor because of accumulated surface error and multiple reflections of the s -polarized beam between plates. The reflected beam is even worse; surface error affects reflected light more than transmitted, and the multiple reflections are not superimposed. The only real advantages are high power handling capability and ready availability of materials.

[†]The word *dichroic* has been given so many different meanings in optics that it's now next to useless.

6.5.2 Multilayer Polarizers

Alternating layers of high and low refractive index can be made into an effective polarizer, similar to the pile of plates but without the beam quality disadvantage. The similarity is not perfect, because interference effects cannot be ignored in thin films, even for white light.

We saw this trick in Section 5.4.4 with polarizing cubes, but such a film stack can also be deposited on a glass plate, forming a polarizing plate beamsplitter. Since there is no optical cement, the damage threshold of these devices is high, making them a good match for powerful pulsed lasers, such as ruby (694 nm) and Nd:YAG (1064 nm). They are rarely used elsewhere, for three main reasons: Brewster incidence is very oblique, so that the reflected light comes off at an odd angle; the angular alignment is critical (as in all Brewster polarizers), and there is no obvious cue for rough alignment as there is in polarizing cubes; and the large index discontinuity at the top surface of the film reflects an appreciable amount of p -polarized light, making the polarization purity of the reflected wave poor.

6.5.3 Polarizing Cubes

Next to film polarizers, the most common type of polarizer in lab drawers is the polarizing beamsplitter cube, which we discussed at length in Sections 4.7.2 and 5.4.5. These are superficially attractive devices that in the author's experience cause more flaky optical behavior than anything else, barring fiber.

6.6 BIREFRINGENT POLARIZERS

Birefringent materials can be used in various ways to make polarizers. The three main classes use (best to worst) double refraction, walkoff, and TIR.

Crystal polarizers are usually made of calcite because of its high birefringence, good optical properties, and reasonable cost. Quartz is sometimes used, but its optical activity causes quartz prisms to exhibit weird polarization shifts versus wavelength, field angle, and orientation—it matters which way round you use a quartz polarizer. None of these crystal devices is cheap, so use them only where you need really good performance. The CVI Laser catalog has an extensive discussion of polarizing prisms.

6.6.1 Walkoff Plates

A very simple polarizer or beam displacer based on beam walkoff (Section 6.3.5) can be made from a thick plate of birefringent material whose optic axis is not parallel to its faces, as in Figure 6.1. When a beam comes in near normal incidence, the o -ray passes through such a plate nearly undeviated, whereas the e -ray walks off sideways. The two are refracted parallel to the original incident beam when they leave the plate. Walkoff plates are inexpensive, because single plane-parallel plates are easy to make, and because no great precision is required in the orientation of the optic axis if only the o -ray is to be kept. This technique is frequently used in optical isolators for fiber applications, where cost is a problem and the angular acceptance is small. Note that the optical path length seen by the two beams is very different, so using walkoff plates as beamsplitters in white-light interferometers is difficult. The shift is a small fraction

of the length of the prism but works over a large range of incidence angles; thus the étendue is small if you need the beams to be spatially separated, but large if overlap is OK.

6.6.2 Savart Plates

The walkoff plate can be made more nearly symmetrical by putting two of them together to make a Savart plate. These consist of two identical flat, square plates of quartz, calcite, or LiNbO_3 whose optic axes are oriented at 45° to the surface normal. The plates are rotated 90° to each other and cemented together. (One plate's optic axis lies in the plane of the top edge and the other one's in the plane of the side edge.)

An o -ray in the first plate turns into an e -ray in the second, and vice versa, so that the two polarizations are offset by the same amount from the axis, in opposite directions, and emerge parallel to their initial propagation direction. At normal incidence, they have zero path difference and hence produce white-light fringes if they overlap.

Away from normal incidence, these are not zero path difference devices, since the e -ray is polarized at a large angle from the optic axis, the path difference changes linearly with angle, rather than quadratically as in a Wollaston prism; Section 19.1.1 has an example where this seemingly obscure point caused a disaster.

6.7 DOUBLE-REFRACTION POLARIZERS

Double-refraction polarizers exploit the different index discontinuity seen by e - and o -rays at an interface. Generally they have excellent performance, but like other refracting prisms their deflection angles change with λ , and they anamorphically distort the beam to some degree.

6.7.1 Wollaston Prisms

A Wollaston prism consists of two wedges of calcite, with their optic axes oriented as shown in Figure 6.2a (the diagram shows a Wollaston prism made from a positive uniaxial crystal such as quartz). A beam entering near normal incidence is undeviated until it encounters the buried interface. There, the e -ray will see the index go down at the surface and so will be refracted away from the normal, whereas the o -ray will see an index increase and be refracted toward the normal by nearly the same amount. (The bending goes the other way for negative uniaxial crystals.) Both beams hit the output facet at a large angle and so are refracted away from the axis.




Figure 6.2. Double-refraction polarizers have the best extinction and purity of any type. (a) Wollaston prisms have no etalon fringes. (b) Rochon prisms have one undeviated beam.

The result is a polarizing prism of unsurpassed performance: polarization ratios of 10^{-6} , nearly symmetrical beam deviation, and, crucially, no internal back-reflection to cause etalon fringes. The beams are not perfectly symmetrical because of the asymmetric incidence on the buried interface. The phase shift between beams is linear in y , independent of x , and varies only quadratically in the field angle, since the optic axis lies in the plane of the prism faces, making Wollastons good for interferometers. You can find a lateral position where the OPD between the two beams is zero, and moving the prism sideways makes a nice phase vernier. Quartz Wollastons have beam separations of 1° to 3.5° , while calcite ones are typically 10° to 20° , and special three-element calcite ones can achieve 30° . Wollastons have excellent étendue on account of their wide angular acceptance.

6.7.2 Rochon Prisms

Table 6.1 shows that the refractive indices of calcite are in the range of optical glass. You can thus make a Wollaston-like prism by using one calcite wedge and one glass one (with $n_{\text{glass}} \approx n_o$), as shown in Figure 6.2. The difference is that one beam is undeviated, as in a Glan–Thompson. Rochon prisms suffer from severe etalon fringes due to the undeviated path, but the polarization purity is similar to a Wollaston, and because of the larger angular acceptance of double refraction, you can use the Rochon tipped fairly far to avoid the back-reflection. The undeviated beam is an o -ray and so has no major chromatic problems.

Some Rochons (and the closely related Senarmont prisms) are made with the glass wedge replaced by a chunk of calcite with its optic axis normal to the output facet, so that both rays see nearly the o -ray index in the second wedge. It has the optical properties of the Rochon without its cost advantage, but is better color-corrected and less likely to delaminate due to differential thermal expansion.

Because of the variations of n_e with incidence angle, a prism (like the Rochon) that transmits the o -ray undeviated is probably superior in imaging applications, as it is easier to correct the resulting aberrations, particularly astigmatism, anamorphic distortion, and chromatic aberration.

Aside: Quartz Rochon Prisms. Occasionally a Rochon prism will surface that has both wedges made of α quartz; unfortunately, optical activity in the output wedge will

TABLE 6.1. Properties of Common Birefringent Materials

Material	n_o	(n_\perp)	n_\parallel	$n_\parallel - n_\perp$	$\lambda(\mu\text{m})$
Magnesium fluoride	1.3783	1.3906	0.0123	540	0.12–8
α Crystal quartz SiO_2	1.5462	1.5554	0.0092	546	0.15–2, 3.3–6
Sapphire Al_2O_3	1.768	1.760	–0.0080	589.3	0.2–6
Potassium dihydrogen phosphate (KDP)	1.5125	1.4705	–0.0420	540	0.2–1.5
Ammonium dihydrogen phosphate (ADP)	1.5274	1.4814	–0.0460	540	0.2–2.0
Barium titanate (BaTiO_3)	2.459	2.400	–0.0590	550	0.4–2
Lithium niobate (LiNbO_3)	2.3165	2.2285	–0.0880	550	0.3–8
Calcite (CaCO_3)	1.658	1.486	–0.1720	589.3	0.25–2.5
Rutile (TiO_2)	2.64	2.94	–0.3000	560	0.4–9




Figure 6.3. Wollaston prisms have a variety of uses: (a) beam splitting, (b) beam combining, (c) heterodyne interferometry, and (d) solid Fourier transform spectrometer.

severely degrade the polarization purity—not only rotating it, but rotating light at different locations and wavelengths by different amounts.

6.7.3 Cobbling Wollastons

Wollaston prisms are usually used as etalon-fringe-free polarizing beamsplitters, but that doesn't exhaust their usefulness. They can of course be run backwards as beam combiners, where two input beams are combined into a single output beam. More generally, a Wollaston can be used to make two beams of orthogonal polarization cross outside the prism, as shown in Figure 6.3.

Due to the slightly unequal angular deviations of the two beams, Wollaston prisms have a path difference that depends on position; an interferometer can have its operating point[†] adjusted by moving the prism back and forth slightly in the direction of the wedges. Note that if you're using the prism to combine beams of any significant NA, this asymmetry makes the plane of the fringes not exactly perpendicular to the axis, so that some in-out motion is needed to maintain the same average optical path in the two beams as you move the prism sideways.

Wollastons are quite dispersive, so they aren't especially useful for imaging in white light unless the beam separation is very small, as in *schlieren interferometry*, which takes advantage of the colored fringes formed by very small-angle Wollastons to make images of weak phase objects such as air currents.

Example 6.1: Solid Fourier Transform Spectrometer. Interferometers based on Wollaston prisms have been used fairly widely. One interesting approach is static Fourier transform interferometry, similar to an FTIR spectrometer (see Section 10.5.6) but with no moving parts. The wide angular acceptance of Wollastons makes high étendue interferometers easy to build. The limiting factor in such interferometers is the birefringence

[†]The operating point of an interferometer, amplifier, or what have you is the nominally stable point of a nonlinear response about which the (supposedly small) signal excursions occur.

of the plates, which makes the transmission phase of an off-axis ray a peculiar function of the angle and limits throughput. Using a positive uniaxial crystal (e.g., quartz) for the splitting Wollaston and a negative uniaxial one (e.g., ammonium dihydrogen phosphate) for the recombiner results in the birefringence canceling out, so that the full étendue is available[†].

6.7.4 Nomarski Wedges

A Nomarski wedge is a modified narrow-angle Wollaston prism. In the Wollaston, the beams diverge from the middle of the prism, so that building an interference microscope requires the microscope objective's exit pupil to be outside the lens, whereas it's usually just inside. The Nomarski prism has the optic axis of one of the wedges oriented out of the plane of the end face; thus the *e*-ray walks off sideways far enough that its exit pupil (where the *e*- and *o*-rays cross and the white-light fringes are located) is outside the prism. In a symmetrical system using two Nomarski wedges, the path difference between the two beams is zero, as is required for achromatic differential interference contrast (DIC) measurements.

6.7.5 Homemade Polarizing Prisms

The one serious drawback to birefringent prisms is that they're expensive, especially in large sizes. You can make an adjustable prism similar to a Wollaston of a few milliradians by bending a bar of plastic such as polycarbonate. Stress birefringence splits the polarizations, and the bending causes the bar to become trapezoidal in cross section (by the local strain times Poisson's ratio) so that the two polarizations are refracted in slightly different directions. This has been used in schlieren interferometers.[‡] The material creeps, so these prisms aren't very stable, and cast plastic isn't too uniform, so they have to be used near an image.

6.8 TIR POLARIZERS

The third major class of birefringent polarizers is based on TIR at a thin layer between two prisms of birefringent material. Because n_e and n_o are different, the *e*- and *o*-ray critical angles will be different as well, so that we can transmit one while totally reflecting the other. In a sufficiently highly birefringent material, the difference is large enough to be useful, although TIR polarizers always have a much narrower angular acceptance than double-refraction ones. The exit angles are set by *o*-ray reflections, so they are pretty well achromatic as long as the exit face is close to perpendicular to the beam (see Section 6.8.1).

The small angular acceptance leads to small étendue, and there are some familiar drawbacks such as poor polarization purity in the reflected beam.[§] A more subtle problem is

[†]D. Steers, B. A. Patterson, W. Sibbett, and M. J. Padgett, Wide field of view, ultracompact static Fourier transform spectrometer. *Rev. Sci. Instrum.* **68**(1), 30–33 (January 1997).

[‡]S. R. Sanderson, *Rev. Sci. Instrum.* **76**, 113703 (2005).

[§]It's a bit more complex than in a polarizing cube, because an oblique reflection at the TIR surface can mix *e*- and *o*-rays. Real prisms are carefully cut to avoid this effect.

that since calcite has a negative birefringence, it's the e -ray that is transmitted undeviated, and in imaging applications, its wavefronts are aberrated by the variation of n_e with angle (see the earlier Rochon prism discussion). All in all, TIR polarizers are inferior to double-refraction types for most uses.

6.8.1 Refraction and Reflection at Birefringent Surfaces

When calculating the behavior of obliquely incident light at the surface of a birefringent material, life gets somewhat exciting unless the optic axis is perpendicular to the plane of incidence. When doing this sort of problem, remember that phase is phase is phase—you calculate phase matching at the interface based on the \mathbf{k} vector of the incoming beam, period. The angle of reflection is a consequence of the fundamental physics involved, namely, the phase matching condition, which remains in force.

For example, imagine making the calcite prism of Figure 6.4a out of MgF_2 , so that the e -ray is now the one reflected. Light coming in s -polarized is a pure o -ray, but the p -polarized light starts out as a high index e -ray and winds up as a low index e -ray (if the wedge angle were 45° it would wind up as an o -ray). Thus the value of k changes, so the “law of reflection” is broken: $\theta_r \neq \theta_i$.

6.8.2 Glan–Taylor

As shown in Figure 6.4a, we can make a TIR polarizer from a simple triangular calcite prism, with the optic axis lying parallel to one edge of the entrance face, and with a wedge angle α whose complement lies between the e - and o -ray critical angles. It has calcite's wide transmission range (220–2300 nm), and because there are no cemented joints, its laser damage threshold is high, $100\text{W}/\text{cm}^2$ or so. This simple device has some serious disadvantages, too; as in the beamsplitter cube, the reflected polarization purity is poor, but there are more. The transmitted beam exits near grazing, because the two refractive indices are not very different; it is anamorphically compressed, which is usually




Figure 6.4. TIR polarizing prisms: (a) simple calcite prism ($n_e < n_o$); (b) quartz Glan–Taylor ($n_e > n_o$) adds an air-spaced second prism to straighten out the transmitted beam; and (c) calcite Glan–Thompson ($n_e < n_o$) uses cement with $n_e < n < n_o$.

undesirable, and it also exhibits chromatic dispersion. The angle of incidence can range only between the e - and o - ray critical angles, which limits the field severely.

If we put two such prisms together, as in Figure 6.4b, we have the Glan–Taylor prism. Here the first prism does the polarizing, and the second one undoes the beam deviation and almost all the chromatic dispersion. The Glan–Taylor keeps the wide transmission and high damage threshold of the calcite prism, but the stray light is quite a bit worse due to multiple bounces in the air gap, where the reflection coefficient is high due to the oblique incidence.

You don't usually try to use the reflected beam from a Glan–Taylor for anything, because its polarization is impure and changes with angle due to the birefringence of the entrance prism. For laser applications, you can cut the prisms at Brewster's angle for the e -ray, which lies above the o -ray critical angle. The resulting Glan-laser prism reduces the angular acceptance while improving the multiple reflections and stray light.

6.8.3 Glan–Thompson

The Glan–Thompson prism of Figure 6.4c is made by cementing two skinny calcite wedges together, hypotenuse to hypotenuse, using cement with an index of about 1.52. The superficial similarity to the Glan–Taylor is somewhat misleading; because n_{glue} is between the e - and o -ray indices, the e -ray cannot experience TIR, so the Glan–Thompson has a much wider angular acceptance than the Glan–Taylor, even though it is longer and skinnier.

As in the Glan–Taylor, the first prism does the polarizing, and the second straightens out the transmitted beam, so that the transmitted beam is undeviated. Because the indices of the calcite and the cement are not very different ($n_o = 1.655$, $n_{\text{glue}} \approx 1.52$), this requires near-grazing incidence, making the Glan–Thompson prism rather long for its aperture.

The o -ray makes two or three TIR bounces in the entrance prism, so that's a good place for a beam dump; Glan–Thompson prisms are usually embedded in hard black wax (a reasonable index match to the o -ray in calcite), so that only the undeviated beam emerges. With a four-sided prism on the entrance side, the reflected ray can be allowed to leave through its own facet, near normal incidence. This configuration is called a beamsplitting Thompson prism and is quite a good device; the closer index match at the interface makes the reflected polarization purer, and the reflected light doesn't see any serious birefringence since its polarization direction is unaltered. Nonetheless, Glan–Thompson prisms share most of the disadvantages of polarizing cubes, including strong etalon fringes and low damage threshold ($\approx 1 \text{ W/cm}^2$) due to the glue, and reflected light polarization purity inferior to that of double-refraction polarizers.

6.9 RETARDERS

The polarization of a monochromatic electromagnetic wave can be decomposed in terms of two arbitrary orthonormal complex basis vectors. This means that, for example, a linearly polarized light beam can be expressed as the sum of two orthogonal circularly polarized beams (of right and left helicity), and vice versa. The devices of this section all use this property to apply different phase delays to different polarization components.

6.9.1 Wave Plates

Retarders or *wave plates* are the simplest application of birefringence. A uniaxial plate of thickness d with its optic axis parallel to its faces will delay a normally incident o -ray by

$$\Delta t_o = n_o \frac{d}{c} \quad (6.8)$$

and similarly for the e -ray, so that the phases of the two are shifted by

$$\delta = (n_e - n_o) \frac{\omega d}{c}. \quad (6.9)$$

Retarders based on this principle are called *wave plates*. When $\Delta\phi$ is $\lambda/4$ for a given wavelength, you have a *quarter-wave plate*, and when it's $\lambda/2$, a *half-wave plate*; these are the two most useful kinds. Note that in the absence of material dispersion, this retardation is a pure time delay, so that the phase varies rapidly with λ . As we saw in Chapter 4, there are also retarders based on the phase shift upon TIR whose phase shift is nearly constant with λ . There are also achromatic wave plates made of multiple polymer layers, whose phase shift is reasonably independent of λ . As usual, there is a three-way trade-off between complexity (i.e., cost and yield), degree of achromatism, and bandwidth.

6.9.2 Quarter-Wave Plates

Quarter-wave plates are good for converting linear to circular polarization and back. Consider a $\lambda/4$ plate lying in the xy plane whose fast axis is \mathbf{x} , with a plane wave passing through it with \mathbf{k} parallel to \mathbf{z} . If \mathbf{E} is at $\pi/4$, the e -ray and o -ray amplitudes will be equal. These get out of phase by $\pi/2$ crossing the plate, so that the field exiting is

$$\mathbf{E}_{\text{out}} = \frac{E_{\text{in}}}{\sqrt{2}} (\hat{\mathbf{x}} \cos \omega t + \hat{\mathbf{y}} \sin \omega t), \quad (6.10)$$

which is left-circular polarization (goes like a right-handed screw). Putting the slow axis along x (or equivalently putting \mathbf{E} at $-\pi/4$) makes right-circular polarization instead.

6.9.3 Half-Wave Plates

A half-wave plate delays one Cartesian component by half a cycle with respect to the other, which reflects \mathbf{E} through the fast axis. This is very useful where both the initial and final polarization states are linear—you twist the wave plate until the polarization lines up just right. Linear polarization stays linear, but with circular or elliptical polarization, the helicity gets changed, so right and left circular are exchanged. Figure 6.5 shows how this works.

Combining differently oriented retarders is most easily done with rotation matrices; a retarder of $\beta\lambda$ whose slow axis is at θ_1 can be written

$$R(\beta, \theta_1) = \begin{bmatrix} \cos \theta_1 & \sin \theta_1 \\ -\sin \theta_1 & \cos \theta_1 \end{bmatrix} \begin{bmatrix} \exp(i2\pi\beta) & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \cos \theta_1 & -\sin \theta_1 \\ \sin \theta_1 & \cos \theta_1 \end{bmatrix}, \quad (6.11)$$

which we'll come back to a bit later.




Figure 6.5. Retarders. Here two quarter-wave plates ($n_{\perp} < n_{\parallel}$) turn a linearly polarized beam first into circular, then into the orthogonal linear polarization, just as a half-wave plate would.

6.9.4 Full-Wave Plates

The retardation of a wave plate depends on its rotations about x and y , so that small errors can be removed by tilting it. For a uniaxial material, the retardation is increased by rotating it around the optic axis, since n_o and n_e are unaltered and the path length is increased; tipping the other way also increases it, because in almost all uniaxial materials, $|n_e - n_o|$ declines more slowly than the path length grows.

A full-wave plate nominally does nothing at all, but in fact the dependence of the retardation on the tipping of the plate makes it a sensitive polarization vernier control, able to provide small amounts ($< \lambda/8$) of retardation to balance out other sources of birefringence in the beam.

6.9.5 Multiorder Wave Plates

Much of the time, it is impractical to make the plate thin enough for such small retardations (a visible quarter-wave plate made of quartz would be only $20 \mu\text{m}$ thick). For narrowband light propagating axially, it is sufficient that the retardation be an odd multiple of $\lambda/2$ for a half-wave plate or $\lambda/4$ for a quarter-wave plate. It is more practical to make a 20.25λ plate (say), which has just the same effect.

Neglecting the change in n_e with angle, the retardation goes approximately as the secant of the incidence angle, so that a 20.25λ multiorder plate will have an étendue over three orders of magnitude smaller than a 0.25λ zero-order one for a given retardation tolerance, and chromatic and temperature shifts 80 times larger as well. This isn't usually a problem with laser beams used in the lab but is serious when you need to run a significant aperture or work in field conditions. One good thing about multiorder wave plates is that they can be made to work at more than one wavelength; for example, you can get quarter-wave plates that work at 633 nm and 1064 nm. It's easier to make them when the wavelengths are significantly different. That kind are often cut with the optic axis slightly out of the plane of the end faces, to make the ratios come out just right with a reasonable thickness. Small amounts of walkoff will result.

6.9.6 Zero-Order Wave Plates

The poor étendue of multiorder wave plates can be fixed by putting two of them together, oriented at 90° to each other so that the aperture-dependent phase shifts largely cancel.

We might laminate a 20.38λ plate with a 19.88λ plate to make a zero-order 0.50λ plate.

6.9.7 Film and Mica

The cheapest retarders are made from very thin layers of mildly birefringent materials such as PVA film or mica. These materials are easily prepared in very thin sheets, and since only a small retardation is desired, only a little is needed. Unfortunately their accuracy, uniformity of retardation across the field, and transmitted wavefront distortion have historically been poor. On the other hand, because even the cheap ones are usually zero-order devices, their dispersion and angular sensitivity are small. Polymer retarders are enjoying a bit of a renaissance, because when carefully made (usually in two-layer stacks) they can have unique properties, such as decent achromatism. In a low coherence imaging system, you can get around their wavefront distortion by putting them near a focus.

6.9.8 Circular Polarizers

The usual polarizer types resolve the incoming polarization state into two orthogonal linear polarizations. This is not of course the only choice; since an arbitrary linear polarization can be built up out of right and left circularly polarized light with a suitable phase shift between them, it follows that we can use these as basis vectors as well. Circular polarizers are quite useful for controlling back-reflections (e.g., glare on glass panels). They're made by laminating a film polarizer to a polymer film quarter-wave plate, with the fast axis of the plate at 45° to the polarization axis. One serious gotcha is that there's usually a retarder on only one side. If used in one direction, that is, with the polarizer turned toward an unpolarized light source, this does result in a circularly polarized output and will indeed attenuate reflected light returned back through it; but it won't pass a circularly polarized beam through unchanged, nor will it work if it's turned round. For that you need two retarders, one on each side.

6.10 POLARIZATION CONTROL

6.10.1 Basis Sets for Fully Polarized Light

We saw in Section 6.2.5 that light could be expressed in terms of Cartesian basis vectors, the *Jones vectors*:

$$\mathbf{E}_\perp = \begin{bmatrix} E_x \\ E_y \end{bmatrix} = E_x \hat{\mathbf{x}} + E_y \hat{\mathbf{y}}. \quad (6.12)$$

A similar decomposition in terms of circularly polarized eigenstates is useful in discussing optical activity and Faraday rotation. A plane wave propagating toward positive Z with complex electric field $\tilde{\mathbf{E}}_\perp$ can be decomposed as

$$\tilde{\mathbf{E}}_\perp \equiv \begin{bmatrix} \tilde{E}_x \\ \tilde{E}_y \end{bmatrix} = \frac{\tilde{E}_L}{\sqrt{2}} \begin{bmatrix} 1 \\ i \end{bmatrix} + \frac{\tilde{E}_R}{\sqrt{2}} \begin{bmatrix} 1 \\ -i \end{bmatrix}, \quad (6.13)$$

TABLE 6.2. Jones Matrix Operators for Common Operations

Coordinate rotation of θ	$\begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix}$
δ Radian retarder, slow axis along \mathbf{x}	$\begin{bmatrix} e^{i\delta/2} & 0 \\ 0 & e^{-i\delta/2} \end{bmatrix}$
Analyzer along \mathbf{x}	$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$
Analyzer at θ to \mathbf{x}	$\begin{bmatrix} \cos^2 \theta & \sin \theta \cos \theta \\ \sin \theta \cos \theta & \cos^2 \theta \end{bmatrix}$

where left- and right-circular components \tilde{E}_L and \tilde{E}_R are given by

$$\tilde{E}_L = \mathbf{E} \cdot \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -i \end{bmatrix} \quad \text{and} \quad \tilde{E}_R = \mathbf{E} \cdot \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ i \end{bmatrix}. \quad (6.14)$$

Linearly polarized light has $\tilde{E}_L = e^{i\phi} \tilde{E}_R$, where ϕ is the azimuthal angle of \mathbf{E} measured from the x axis. Polarization gizmos can be represented as 2×2 matrix operators; in Cartesian coordinates, the result is the *Jones matrices* shown in Table 6.2.

Like *ABCD* matrices, these are not powerful enough to model everything without hand work; for example, treating reflections takes some care. We model optical activity and Faraday rotation as coordinate rotations, but since one adds and the other cancels on the way back, we have to be very careful about the bookkeeping; we'll do an example of this later when we consider the Faraday rotator mirror.

6.10.2 Partial Polarization and the Jones Matrix Calculus

Light can be partially polarized as well, most often by reflecting thermal light from a dielectric. This polarization is often fairly strong, especially when the reflection takes place near Brewster's angle, as we know from the effectiveness of Polaroid sunglasses in reducing glare from water, car paint, and glass. Rayleigh scattering also produces partially polarized light; try your sunglasses on a clear sky—when you look at 90° to the sun, the polarization is quite pronounced.

The vectors that are adequate for describing full polarization fail for partial polarization, which is an intrinsically more complicated situation. If the direction of \mathbf{k} is prespecified, we can express the polarization properties of a general narrowband beam as a four-dimensional vector (the *Stokes parameters*, see Born and Wolf) or by a 2×2 *coherency matrix*.[†] As Goodman explains,[‡] the coherency matrix formulation lets us follow the polarization state of our beam through the system by matrix multiplication of \mathbf{E}_\perp by the accumulated operator matrices, written in reverse order, just the way we did ray

[†]The two are very closely related; the elements of the coherency matrix are linear combinations of the Stokes parameters.

[‡]Joseph W. Goodman, *Statistical Optics*. Wiley, Hoboken, NJ, 1986.

TABLE 6.3. Coherency Matrices for Some Polarization States

Linear x	$E_0 \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$	Linear y	$E_0 \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$
Right circular	$E_0 \begin{bmatrix} 1 & i \\ -i & 1 \end{bmatrix}$	Left circular	$E_{0/2} \begin{bmatrix} 1 & -i \\ i & 1 \end{bmatrix}$
Unpolarized	$E_{0/2} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$		

tracing with the $ABCD$ matrices, and it is easily connected to time-averaged polarization measurements. The coherency matrix $\underline{\mathbf{J}}$ is the time-averaged direct product $\underline{\mathbf{E}}\underline{\mathbf{E}}^{\dagger}$:

$$\underline{\mathbf{J}} = \begin{bmatrix} J_{xx} & J_{xy} \\ J_{xy}^* & J_{yy} \end{bmatrix} = \begin{bmatrix} \langle \tilde{E}_x \tilde{E}_x^* \rangle & \langle \tilde{E}_x \tilde{E}_y^* \rangle \\ \langle \tilde{E}_y \tilde{E}_x^* \rangle & \langle \tilde{E}_y \tilde{E}_y^* \rangle \end{bmatrix}. \tag{6.15}$$

It’s easy to see from the definition that (up to a constant factor) J_{xx} and J_{yy} are the real-valued flux densities you’d measure with an analyzer along $x(I_{0^\circ})$ and $y(I_{90^\circ})$, respectively. The complex-valued J_{xy} is related to the flux density I_{45° that you get with the analyzer at 45° , and the I'_{45° you get by adding a $\lambda/4$ plate with its slow axis along y before the analyzer,

$$J_{xy} = I_{45^\circ} - \frac{1}{2}(I_{0^\circ} + I_{90^\circ}) + i \left(I'_{45^\circ} - \frac{1}{2}(I'_{0^\circ} + I'_{90^\circ}) \right) \tag{6.16}$$

Table 6.3 has coherency matrices for some special cases.

6.10.3 Polarization States

It is commonly held that when you superpose two beams, their $\underline{\mathbf{J}}$ s add, but that assumes that they are mutually incoherent, which is far from universally true. You’re much safer sticking closer to the fields unless you know *a priori* that the waves you’re combining are mutually incoherent but still narrowband enough for the Jones matrix approach to work.

A couple of mathematical reminders: a lossless operator $\underline{\mathbf{L}}$ is unitary—all its eigenvalues are on the unit circle and $\underline{\mathbf{L}}\underline{\mathbf{L}}^\dagger = \underline{\mathbf{I}}$, that is, $\underline{\mathbf{L}}^\dagger = \underline{\mathbf{L}}^{-1}$, where the adjoint matrix $\underline{\mathbf{L}}^\dagger$ is the complex conjugate of the transpose, $\underline{\mathbf{L}}^\dagger = (\underline{\mathbf{L}}^T)^*$. These lists can be extended straightforwardly by using the definition (6.15) and matrix multiplication. Remember that although the operators multiply the *fields* (6.13) directly, applying a transformation to $\underline{\mathbf{J}}$ or $\underline{\mathbf{L}}$ requires applying it from both sides; if $\mathbf{E}'_\perp = \underline{\mathbf{L}}\mathbf{E}_\perp$,

$$\langle \mathbf{E}'(\mathbf{E}')^* \rangle^T = (\underline{\mathbf{L}}\mathbf{E})(\underline{\mathbf{L}}^*\mathbf{E}^*)^T = \underline{\mathbf{L}}\underline{\mathbf{J}}\underline{\mathbf{L}}^\dagger. \tag{6.17}$$

It’s worth writing it out with explicit indices and summations a few times, if you’re rusty—misplacing a dagger or commuting a couple of matrices somewhere will lead you down the garden path otherwise.

6.10.4 Polarization Compensators

Extending the zero-order quartz wave plate idea, we can make one plate variable in thickness by making it from two very narrow wedges rather than a single plate, yielding the *Soleil compensator*. Provided the two wedges are identical, the retardation is constant across the field and can be adjusted around exactly 0 by sliding one of them, which is a great help in wide-field and wideband applications, for example, looking at small amounts of stress birefringence with a nulling technique. Exactly zero retardation is important only in such cases; in narrowband, low NA systems, it's enough to have the retardation be $0 \bmod 2\pi$, and these compensators are no better than a pair of quarter-wave plates.

6.10.5 Circular Polarizing Film for Glare Control

Laminated circular polarizers of moderate performance can be made cheaply in very large sizes, which makes them attractive as a glare-reduction device for instrument displays; room light passing through the polarizer and then being reflected is absorbed on the return trip. Before CRT screens were AR coated efficiently, circular polarizers were very popular computer accessories, and they remain useful in similar situations.

6.10.6 Polarization Rotators

Optically active materials such as quartz or sugar solutions can be used for polarization rotators. Those made from amorphous materials (e.g., Karo corn syrup) have no birefringence, and so linearly polarized light stays linearly polarized regardless of wavelength, field angle, or what have you. They're inconvenient to make, hard to adjust, and ten times more dispersive than half-wave plates, so apart from special situations such as optical diodes, they have few compelling advantages.

6.10.7 Depolarizers

It is impossible to reproduce the extremely rapid polarization changes of thermal light when starting from a more coherent source such as a laser beam. Devices that purport to depolarize light never do it well enough for that; they just produce a speckle pattern varying more or less rapidly in time or space. If your experiment is slow enough, this may suffice, but in fact it rarely does. Polarized light is something we just have to live with.

There are two classes of depolarizers: wave plates whose retardation varies rapidly across their faces (e.g., Cornu depolarizers), and moving diffusers, such as a disc of ground glass spun rapidly on a shaft. A Cornu depolarizer can do a reasonable job on wideband light, providing the 2π period of the polarization change is sufficiently rapid and the spatial resolution sufficiently low.

Fixed depolarizers help to eliminate the mild polarization dependence of some optical instruments, for example, PMTs, grating spectrometers, and so on, when used with broadband light that may be partially polarized. They do a good enough job for that, certainly.

The spinning ground glass technique often tried with laser beams is much less successful: all you get is a rapidly rotating speckle pattern, which causes a whole lot of

noise. Unlike the situation in Section 2.5.3, the order-1 changes in instantaneous intensity at any point are not smeared out over a bandwidth of hundreds of terahertz, but concentrated in a few hundred kilohertz; the result is not pretty. The rotating speckle pattern also produces undesired correlations in space. These correlations can be reduced by using two discs rotating in opposite directions; if this is properly done, the speckles no longer appear to rotate. Doing it really properly is not trivial, however, and anyway the huge intensity noise remains. If you are forced to do your measurement this way, be prepared to integrate for a long time; it is normally preferable to use polarization flipping, where you do separate measurements in horizontal and vertical polarization and combine the two, or use a rotating $\lambda/2$ plate and integrate for a whole number of periods.

6.10.8 Faraday Rotators and Optical Isolators

Faraday rotators are straightforward applications of the Faraday effect, using a magneto-optic crystal such as YIG in a magnetically shielded enclosure with carefully placed, stable permanent magnets inside providing a purely axial field in the crystal. The main uses of Faraday rotators are in optical isolators and in transmit/receive duplexing when the returned light is not circularly polarized, for one reason or another.

These devices, shown in Figure 6.6, all rely on nonreciprocal polarization rotation. The simple isolator uses two walkoff plate polarizers, oriented at 45° to one another, and a 45° Faraday rotator. Light making it through the first rotator gets its \mathbf{E} rotated through 45° on the first pass, so that it is properly oriented to pass through the second polarizer without loss. Light coming the other way has to make it through the second polarizer and is then rotated 45° in the same direction, putting it at 90° to the first polarizer, so that none gets through. Ideally the isolation would be perfect, but it is more typically 30 dB per isolator, with a loss of about 1 dB.




Figure 6.6. Two polarizers plus a 45° Faraday rotator make an optical isolator.

This simplest Faraday isolator requires fully polarized input light, but polarization-insensitive ones can also be made; since you have to use two different paths, it isn't trivial to preserve the input polarization in the process, unfortunately.

The most important uses of Faraday isolators are preventing feedback-induced instability in diode lasers and in preventing high-finesse Fabry–Perot cavities from pulling the frequency of sources being analyzed as the F-P scans.

It is possible to build a *circulator*, an M -port device where the input on port m goes out port $m + 1 \pmod{M}$. Circulators are common in microwave applications but rare in optics. A related device is the *optical diode*, a 45° Faraday rotator plus a -45° optically active cell, used in ring laser cavities; only one propagation direction is an eigenstate of polarization, so that the ring lases in one direction only (the other one gets killed by the Brewster windows).

6.10.9 Beam Separators

A polarizing beamsplitter such as a cube or Wollaston, plus a $\lambda/4$ plate, makes a beam separator, very useful for separating the transmit and receive beams in an interferometer or scanning system. The wave plate is aligned at 45° to the axes of the beamsplitter, as shown in Figure 6.7. On the transmit side, p -polarized light passes through the prism and gets turned into left-circular polarization. Specularly reflected light comes back right-circular, and so the $\lambda/4$ plate turns it into s -polarized light in the cube, which is reflected. If the components are lossless, the wave plate accurate, and the incoming light perfectly p -polarized, the beam suffers no loss whatever in its round trip.

6.10.10 Lossless Interferometers

In Section 4.8.1, we saw that sending light on two passes through a nonpolarizing beamsplitter costs you a minimum of 75% of your light. That's 93.75% of your detected electrical power, representing an SNR degradation of 6 dB in the shot noise limit and 12 dB in the Johnson noise limit—and that's in the best case, with a beamsplitter without excess loss.




Figure 6.7. A polarizing beamsplitter plus a quarter-wave plate make a beam separator, able to disentangle the transmit and receive beams of an interferometer or scanning system.

If we have a polarized source such as a laser, we can use a beam separator to split and recombine the light as shown in Figure 1.12. The only problem is that the two beams are orthogonally polarized and so don't interfere. The usual solution to this is to put an analyzer at 45° to the two beams, resulting in 100% interference but 6 dB detected signal loss. However, by using a polarizing beamsplitter oriented at 45° , detecting the two pairs of beams separately, and subtracting the resulting photocurrents, we get the equivalent of 100% interference with no signal loss, as in the ISICL sensor of Example 1.12.

6.10.11 Faraday Rotator Mirrors and Polarization Insensitivity

As an application of the Jones matrix calculus, let's look at the Faraday rotator mirror, which is widely used in fiber optics. It consists of a 45° Faraday rotator in front of a mirror, so that the light passes twice through the rotator, and so the total rotation is 90° . We have to complex-conjugate the beam to represent the mirror, because the helicity changes and in this model we've no way of expressing the propagation direction.

In operator representation, this is

$$\begin{aligned} \mathbf{E}' &= \mathbf{R}_{\pi/4} (\mathbf{R}_{\pi/4} \mathbf{E})^* = \mathbf{R}_{\pi/2} \mathbf{E}^* \\ \Rightarrow \mathbf{E}'^* \cdot \mathbf{E} &= \begin{bmatrix} E_x & E_y \end{bmatrix} \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \begin{bmatrix} E_x \\ E_y \end{bmatrix} = 0, \end{aligned} \quad (6.18)$$

that is, the light coming back is polarized orthogonally to the light coming in, *regardless of the incoming polarization*—it works for circular and elliptical, as well as linear. It's obviously orthogonal if the polarization exiting the fiber is linear (courtesy of the Faraday rotation) or circular (courtesy of the mirror). Elliptical polarizations have their helicity inverted by the mirror, and their major axis rotated 90° by the Faraday rotator.

The power of this is that the polarization funnies encountered by the beam, that is, birefringence and optical activity, are all unitary operations, so the incoming and outgoing polarizations remain orthogonal everywhere, as long as they traverse the same path. That means that our optical fiber can misbehave as much as it wants, in theory, and as long as we've got a Faraday mirror at the other end, the round-trip light comes out polarized orthogonally to the incoming light; if we send in vertically polarized light, it comes out horizontally polarized, no matter how many waves of birefringence it encountered. This doesn't work quite as well as we'd like, because the accuracy requirements are very high and it ignores scattering, multiple reflections, and transients. Nonetheless, we can build more-or-less polarization-insensitive fiber interferometers this way.

A slightly more subtle benefit is that the propagation phase is polarization insensitive. A lossless fiber has two orthogonal eigenmodes. If we decompose any incoming polarization into these modes, we find that the Faraday rotator mirror exchanges the components in the two eigenmodes, so that the total round-trip phase delay is the sum of the one-way delays of the eigenmodes. You do have to think about Pancharatnam's phase, though (see Section 6.2.4).