

Fiber Optics

All mankind in our age have split up into units, they all keep apart, each in his own groove; each one holds aloof, hides himself, hides what he has from the rest, and he ends by being repelled by others and repelling them.

—Elder Zosima, in *The Brothers Karamazov* by Fyodor Dostoevsky

8.1 INTRODUCTION

We all know that fiber optics and lasers have revolutionized long-haul and high bandwidth terrestrial communications. Fibers are unequaled at that job, there's no doubt about it. They're also good for instruments, though the situation there is less clear.

Fiber optics really constitutes a different domain, a bit like **k**-space, where different things are easier or harder than in bulk optics. There is some new physics available in fibers, and a lot of old physics under a new aspect and in new operating regimes. The trade-offs are different enough that they're worth a considerable amount of study, so we'll concentrate on the characteristics of fibers in instruments. There are a lot of books available on fiber communications and fiber sensors, but not a lot on why to choose fiber or bulk for a particular application.

The thing to remember about using fiber is that it's far harder than it looks. It's seductive; a fiber sensor can easily be made to work at some level, which makes it easy to suppose that making that sensor robust enough for field use is straightforward—and it isn't straightforward at all.

8.2 FIBER CHARACTERISTICS

An optical fiber is a thin cylinder of doped fused silica, with nonuniform n near the center, covered with a protective plastic or metal jacket. It works as a very low loss dielectric waveguide. The most common type is *step-index fiber*, where a central *core* has a uniform index n_1 , which abruptly changes to n_2 where the annular *cladding* begins. The cladding index is also constant out to the edge of the jacket. Most fibers use a germania (GeO_2) doped silica core and a pure silica cladding. (See Table 8.1.)

TABLE 8.1. Fiber Parameters

Core radius	a	2.5–5 μm (single mode) 25–900 μm (multimode)
Core index	n_1	1.47–1.52
Cladding index	n_2	1.464
Normalized index difference	$\Delta = \frac{n_1^2 - n_2^2}{2n_1^2}$	0.004–0.04 (single mode) Δ usually small)
Mode volume	$V = an_1k_0\sqrt{2\Delta}$	1.8–2.4 (single mode)
Fracture strength	S	0.01–4 GPa
NA (multimode)	$\text{NA} = \sqrt{n_1^2 - n_2^2} = n_1\sqrt{2\Delta}$	0.1–0.4
Étendue	$n^2A\Omega$	10^{-8} to 10^{-5} $\text{cm}^2\cdot\text{sr}$

8.2.1 Fiber Virtues

Fibers have lots of virtues, even besides those that all optical components share, such as wide temporal bandwidths.

Cheapness. It's nearly free (10 cents per meter in large quantity for single-mode communications fiber), and components are getting cheaper. Fiber connectors are amazingly cheap for such precise artifacts; they're only a few dollars, as long as you're using 125 μm OD fiber. You can thus bolt stuff together with fibers very quickly, which in principle would allow checking out a measurement idea faster than with bulk optics, where everything has to be aligned manually. And components such as directional couplers are mass produced, which makes them far cheaper than previously.

By putting the optical complexity at the far end, the business end of a fiber sensor can be made very small and cheap, even cheap enough to be disposable, as in fiber catheters, or sacrificial, as in fiber strain gauges cast into concrete structures or shrapnel velocity sensors for ordnance development.

Good Transmission Properties. For long distance transmission, fibers are tops, as already mentioned. They have very low loss (0.2–10 dB/km), exhibit no crosstalk, and transmit over a reasonably wide wavelength range—250–2000 nm at least, and special fibers go somewhat beyond both extremes.

EMI Immunity. Not being inductive or capacitive at ordinary frequencies, fibers are totally immune to electromagnetic interference (EMI), ground loops, and pickup; you can send a very weak signal a long way in a hostile environment. This is a key virtue for distributed sensing and smart structures.

Versatility. Most things that you can do with bulk optics, you can do with fiber, at least in principle. It provides some altogether new capabilities and more convenient ways of doing the same things, or is lower in cost because of economies of scale (driven by fiber transmission), simplified mounting, and reduced size and mass.

Ease of Use. Fibers are mechanically flexible and can bend light around in a pretzel shape if you want. They couple well to many kinds of lasers. Once you've got the light going in and out the ends, no more alignment is necessary. In connectorized systems,

the alignment is done by the manufacturer of the connector. Convenient collimators are available inexpensively.

Environmental Robustness. Fibers, being made almost entirely of quartz, survive high and low temperatures extremely well. When properly jacketed, they are immune to dirt and most fluids, too, though they become much more vulnerable to breakage when wet.

New Physics. The long interaction lengths possible with fiber allow us to base good measurements on intrinsically very weak effects, such as the Faraday effect in silica. We can make all-fiber sensors of completely new types, for example, multiplexed or distributed sensors like OTDR and fiber Bragg grating types.

8.2.2 Ideal Properties of Fiber

Besides these comparative advantages, fiber systems have a few *perfect* characteristics, which can be used to good effect.

Pointing Stability. Single-mode fibers have essentially perfect pointing stability, because with due care, we can ensure that really only one mode reaches the end facet. The mode pattern depends only on wavelength and material properties, so it is extremely stable, merely changing very slightly in diameter and NA with temperature, due to $\partial n / \partial T$ (see Example 9.7).

There-and-Back Phase Symmetry. There is a broad class of cases where the one-way phase delay is exactly the same in both directions in a fiber. The mode structure doesn't change when we replace k_z with $-k_z$, so the propagation velocity of a single mode is unaffected (Figure 8.1). In the face of mode coupling, we have to be a bit more careful. For the mode coupling to undo itself, the fiber has to be lossless and the




Figure 8.1. There-and-back polarization stability with Faraday rotator mirrors.

polarization has to come back in complex conjugated (i.e., with the same major axis and helicity) from the way it went out (see Section 8.3.3 below).

Scatter will degrade the matching stochastically, and there are also three effects that modify it deterministically: Pancharatnam's topological phase (see Section 6.2.4) in non-planar systems, the Sagnac effect in accelerated or rotating ones, and the Faraday effect in magnetically excited ones.

Transient effects will of course arrive at different times in the two directions, depending on where in the fiber they occur. Although silica is a very linear material, in a long fiber path you may see nonlinear effects, for example, the optical Kerr effect that causes phase offsets in Sagnac interferometers.

Two-Pass Polarization Stability. Polarization has two degrees of freedom per mode. Light in a lossless single-mode fiber can thus be resolved into two orthogonal polarization states, and given any one state ϕ_1 , we can always construct a ϕ_2 that is orthogonal to it.

Light launched into the fiber gets its polarization modified wholesale, due to bend birefringence, topological phase, twist-induced optical activity, and other effects, but throughout the process, ϕ_1 and ϕ_2 remain orthogonal, and so maintain their identity if not their appearance. This is a consequence of the unitarity of the Jones matrix of a lossless polarization element (see Section 6.10.2); an infinitesimal length of fiber can be represented by a differential Jones matrix, and the whole fiber's Jones matrix follows by passing to the limit of many segments of very small length.

A Faraday rotator mirror (or FRM, see Section 6.10.11) turns any monochromatic polarization state into its orthogonal partner, which retraces its path through the fiber but remains orthogonal to the first pass. Thus the light reemerging from the input facet has a stable polarization, which is orthogonal to the input light, and so can be separated from it with a polarizing prism. Of course, really high polarization selectivity (e.g., 50 dB optical) requires unfeasibly accurate fiber couplers, end faces, and FRMs.

8.2.3 Fiber Vices

Fiber has its own peculiar vices as well, some of which are crippling, and few of which are obvious. Most stem from having a too accurately aligned system with a long path length in glass that's hard to get light into. Attentive readers may note a certain similarity between the lists of virtues and vices—that's where our discussion of fibers in instruments centers.

Fiber Isn't So Cheap in Real Life. That 10 cents per meter pricing requires you to buy many kilometers of fiber, and it only applies to garden-variety single-mode telecom fiber—not PM or multimode fiber. Fiber components and connectors are also not cheap except at 1310 and 1550 nm and 125 μm OD.

Mechanical Fragility. It's made of glass, after all, and is vulnerable to snagging, kinking, and abrasion. Bending also causes loss.

Really Tiny Étendue. It's so low that fiber doesn't couple well to anything but focused laser beams. The same effect is responsible for fiber's horrible etalon fringe problem—there's nowhere for reflections to go except exactly back the way they came. (Angle polishing helps a lot, but not enough.) Its power handling capacity is limited, too.

Mode Instability. Multimode fiber, especially if it has only a few modes, exhibits large amounts of noise and drift in its illumination pattern due to mode coupling.

Etaion Fringes and FM–AM Conversion. Every optical system is an interferometer, and fiber optic systems are several interferometers at once, each one demodulating phase noise like a delay line discriminator. The too accurate alignment leads to severe etalon fringes and laser feedback and (as we saw in Section 2.5.3) removes the spatial averaging that helps so much to eliminate the FM–AM conversion from interference with delayed light.

Besides the obvious facet reflections, there are distributed effects that are often even larger; multiple Rayleigh scattering within a long fiber causes serious FM–AM conversion too (see Section 13.5.6). This sounds like a trivial effect, but unfortunately it is not; it is usually by far the dominant broadband AM noise source in a single-mode fiber sensor when the fiber length is a significant fraction of both a coherence length and an absorption length.[†] The physics is similar to the coherence fluctuation noise of Sections 2.5.3 and 19.1.1.

Phase and Polarization Instability. The long path in glass makes fibers vulnerable to piezo-optic shifts, temperature changes, vibration, and bending, so that every fiber device has serious instabilities of phase and polarization with time.

Sensitivity to Everything Except EMI. Use of fibers for DC measurements is extremely difficult, because apart from pure intensity instruments, every fiber sensor is also a thermometer. For example, fiber strain sensors are plagued by temperature sensitivity, because 1 °C of temperature change gives the same signal as a fairly large strain (10 $\mu\epsilon$), and separating the two effects well enough to get high sensitivity thus requires very accurate, independent temperature measurements, and an accurate knowledge of the response of the system to temperature alone.

8.3 FIBER THEORY

In order to be able to use the good properties of fiber, and avoid being blown up by the bad ones, we need to spend some time on their theoretical properties.

8.3.1 Modes

A step-index fiber can be modeled as a right-circular cylinder of radius a and index n_1 , surrounded by a practically infinite thickness of n_2 . Most of us probably remember how to solve the wave equation in a uniform medium with cylindrically symmetric boundary conditions: the Laplacian operator separates, so we get a one-dimensional wave equation in the axial direction, Bessel's equation in the radial coordinate, and an eigenvalue equation for the azimuthal coordinate, leading to solutions of the form

$$\begin{aligned}\psi_m^{(1)}(r, \phi, z) &= J_m(\kappa r) e^{i(m\phi + k_z z)}, \\ \psi_m^{(2)}(r, \phi, z) &= N_m(\kappa r) e^{i(m\phi + k_z z)},\end{aligned}\tag{8.1}$$

[†]A. Yariv, et al., *IEEE Trans. Lightwave Technol.* **10**(7), 978–981 (1992).

where J_n and N_n are the Bessel functions of the first and second kind, respectively, and

$$k_z^2 + \kappa^2 = k^2 = n_1^2 k_0^2. \quad (8.2)$$

From experience, we expect to apply boundary conditions and get a denumerable set of modes for any given k . For a finite radius, only a finite number can propagate, as usual.

For step-index fibers, we have two dielectrics, so we have to patch the fields together at the boundary, using the phase matching condition in z (because of translational invariance) and continuity in r and ϕ of tangential E and H and of perpendicular D and B , leading to an eigenvalue equation for the allowed values of κ . This requires a significant amount of bookkeeping and results in four families of modes— TE , TM , EH , and HE . For large diameters and small index differences, these can be combined into nearly transverse electromagnetic (TEM) forms, and the scalar approximation is valid; thus we use the scalar patching conditions, that is, phase matching in z and continuity of ψ and $\partial\psi/\partial r$. With a little bit of hand work, this yields the linearly polarized LP modes. While these solutions are approximate, it's much better in practical work to have an intuitive approximation than an unintelligible exact result.

We normally confine our attention to those modes that are guided, so that the field decays with r outside the core of the fiber. Such modes occur when k_z is large enough so that (8.2) forces κ to be imaginary in the cladding while being real in the core, which occurs when

$$n_2 k_0 < k_z < n_1 k_0. \quad (8.3)$$

When κ is imaginary, the solutions are more conveniently written in terms of the modified Bessel functions,

$$\begin{aligned} \psi_m^{(3)}(r, \phi, z) &= I_m(\gamma r) e^{i(m\phi + k_z z)}, \\ \psi_m^{(4)}(r, \phi, z) &= K_m(\gamma r) e^{i(m\phi + k_z z)}, \end{aligned} \quad (8.4)$$

where $k_z^2 - \gamma^2 = n_2^2 k_0^2$. These functions are both monotonic: $I_m(r)$ grows exponentially and is regular at 0, while $K_m(r)$ decays exponentially but has a pole at 0. Because the cladding is very thick and doesn't extend to 0, only the $K_m(\gamma r)$ solution is physical, and the patching conditions lead to an eigenvalue equation for κ , which sets the allowed modes.

This may seem a bit *recherché*, but it has a very important consequence: because we have to patch J_m and $d/dr(J_m)$ together with a monotonically decaying function, $|J_m(\kappa r)|$ must be decreasing at $r = a$. Thus the m th radial mode can only propagate when $d/dr(|J_m(\kappa a)|) \leq 0$. Below the first zero of $J_1'(\kappa a)$, only a single mode can propagate, so the fiber is single mode when

$$\lambda_c > \frac{2\pi}{2.405} a n_1 \sqrt{2\Delta}. \quad (8.5)$$

The light is not confined entirely to the core, as it is in a metal waveguide, but spreads a considerable way into the cladding; the patching condition can always be met for the

lowest mode, so there is no long-wavelength mode cutoff of a fiber the way there is for a metal waveguide. On the other hand, the guiding becomes weaker and weaker as more and more of the light is in the cladding, so there is a practical minimum for the core diameter. Guiding behavior is captured by the *mode volume* or *normalized frequency* V , which from Table 8.1 is

$$V = an_1k_0\sqrt{2\Delta}. \quad (8.6)$$

Good guiding and reliable single-mode operation occurs for V from about 1.8 to 2.4. The lowest order mode, LP_{01} , has a round top and exponentially decaying sides, so it looks roughly Gaussian. The *mode field radius* w_0 is defined similarly, as the $1/e^2$ radius of the best-fit Gaussian, and for single-mode fiber is approximately

$$w_0 \approx a \left(\frac{1}{3} + \sqrt{\frac{2.6^3}{V}} \right). \quad (8.7)$$

As with Gaussian beams, we normally care about the details of only the lowest order mode, because any device that relies on a multimode fiber having exactly predictable mode properties is doomed; a multimode fiber has such strong coupling between modes that the mode amplitudes and phases change rapidly with propagation distance and environmental perturbations. One exception is when we know we have exactly two well-behaved modes, for example, in a polarization preserving fiber (see Section 8.4.4).

8.3.2 Degeneracy

When we talk about a fiber being *single mode*, we are of course talking about the scalar approximation, whereas light in the fiber has two degrees of freedom in polarization. For perfect cylindrical symmetry and homogeneous, isotropic materials, the fiber has rotational symmetry, which means that the two polarization states have exactly the same k_z ; that is, they are *degenerate*. Degenerate states don't get out of phase with each other, so even a very small coupling causes complete energy redistribution between them, as we'll see.

8.3.3 Mode Coupling

Many fiber devices are based on weak lossless (*adiabatic*) coupling between two waveguide modes, either the LP_{01} modes in two fibers brought together, as in a directional coupler, or between two modes in a single fiber. By *weak*, we mean that the mode amplitudes change on a scale much longer than a wavelength, not that the energy transfer isn't large eventually. Under these conditions, we can treat mode coupling by first-order perturbation theory, which is a good break because for that we need only the zero-order waveguide modes. Consider the situation of Figure 8.2: two modes, ψ_1 and ψ_2 , that are orthonormal[†] except for a small coupling, which is constant in z and t . We write the

[†]That is, orthogonal and with the same normalization, so that unit amplitude in one mode corresponds to unit power.




Figure 8.2. Coupled modes.

mode amplitudes as $(A_1, A_2)^T$. The first-order perturbation equation for \mathbf{A} is then

$$\frac{d}{dz}\mathbf{A} = \begin{bmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \end{bmatrix} \begin{bmatrix} A_1 \\ A_2 \end{bmatrix}, \quad (8.8)$$

with $|c_{12}|, |c_{21}| \ll |k_1|, |k_2|$. If the perturbation is lossless, then $d/dz|\mathbf{A}|^2 = 0$. Rearranging a few indices leads to

$$\underline{\mathbf{c}} + (\underline{\mathbf{c}}^T)^* = 0 \Rightarrow \underline{\mathbf{c}} = \begin{bmatrix} ik_{z1} & c_{12} \\ -c_{12}^* & ik_{z2} \end{bmatrix}, \quad (8.9)$$

that is, $\underline{\mathbf{c}}$ is skew-hermitian,[†] and all its eigenvalues are therefore imaginary. We'd expect the pure imaginary elements on the main diagonal, because that's the behavior of the uncoupled modes; the skew-hermitian symmetry of the off-diagonal elements makes sure that the power leaving one mode winds up in the other.

This constant-coefficient system is solved by finding the eigenvalues $i\beta_j$ and eigenvectors ϕ_j of $\underline{\mathbf{c}}$, because in the eigenvector basis (the principal components basis), the equations decouple, leading to solutions $\psi_j(z)$ with a simple exponential z dependence, $\phi_j = \exp(i\beta_j z)$,

$$\beta_j = \frac{1}{2} \left(k_{z1} + k_{z2} \pm \sqrt{(k_{z1} - k_{z2})^2 + 4|c_{12}|^2} \right). \quad (8.10)$$

We'll use primed quantities in this basis (e.g., \mathbf{A}' is the principal component representation of \mathbf{A} , but it's the same vector really). If the two waveguide modes have the

[†]A skew-hermitian or antihermitian matrix is i times a hermitian one, so with that modification we can apply all the usual theorems about hermitian matrices—orthogonal eigenvectors, real eigenvalues, and so on.

same k_z (e.g., the fiber coupler case, or polarization coupling in an ordinary single-mode fiber),

$$\beta_i = k_z \pm |c_{12}| \quad (8.11)$$

and ϕ_j becomes

$$\phi_1 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ i \end{bmatrix}, \quad \phi_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -i \end{bmatrix}. \quad (8.12)$$

This isn't too spectacular looking—just two quadrature pairs with slightly different propagation constants. Now, however, let's consider the usual case, where all the light is initially in $\psi_1 = 1/\sqrt{2}(\phi_1 + \phi_2)$, so $\mathbf{A}(0) = (1, 0)^T$ and $\mathbf{A}'(0) = (\beta_1/\sqrt{2}, \beta_2/\sqrt{2})^T$. In that case, we get a time dependence

$$\mathbf{A}(z) = e^{ik_z z} \begin{bmatrix} \cos(c_{12}z) \\ i \sin(c_{12}z) \end{bmatrix}, \quad (8.13)$$

so that after a distance $\pi/(2c_{12})$, all the energy has moved from ψ_1 to ψ_2 , and it then transfers back and forth sinusoidally, forever. Thus regardless of the magnitude of c_{12} , we can pick a length for our coupler that will give us any desired output power ratio from 0 to ∞ .

If $k_{z1} \neq k_{z2}$, we still get sines and cosines, but with offsets due to the maximum coupling being less than 1; the power doesn't finish moving from ψ_1 to ψ_2 before the phase shift between modes exceeds π , and it starts moving back again. It's useful to define the *beat length* λ_b , which is the period of the spatial beat between modes,

$$\lambda_b = \frac{2\pi}{|k_{z1} - k_{z2}|}. \quad (8.14)$$

Looking at the square root in (8.10), there are two distinct limits,

$$|c_{12}|^2 \gg (k_{z1} - k_{z2})^2/4 \quad (8.15)$$

(long beat length) and

$$|c_{12}|^2 \ll (k_{z1} - k_{z2})^2/4 \quad (8.16)$$

(short beat length), which govern how we do the binomial expansion of the square root (see the problems in the Supplementary Material). In the long beat length limit, the situation is very similar to the degenerate case, except that the maximum coupling is reduced slightly. In the short beat length limit, though, the rapidly growing phase shift between modes prevents much power from coupling at all, so the modes remain more or less uncoupled.

8.3.4 Space-Variant Coupling

The situation is quite different when the coupling coefficient c_{12} is a function of z , especially if we make it periodic near λ_b . If we do so, for example, by launching an acoustic wave of that wavelength, making a photorefractive Bragg grating, or by pressing a corrugated plate against the fiber, we can recover the strong coupling that we had lost. If the coupling has spatial frequency $2k_z$, light going to positive Z will get backscattered strongly, as in a fiber Bragg grating. The coupled-mode theory becomes a bit more complicated in this case, but is still quite manageable.

8.3.5 Dispersion

Because different order modes have different k_z , they have different phase velocities,

$$v_p = \frac{\omega}{k_z} \quad (8.17)$$

and because of the nonlinearity of the expression for k_z in terms of k and κ , their group velocities

$$v_g = \frac{\partial \omega}{\partial k_z} \quad (8.18)$$

are different as well. This *intermodal dispersion*[†] severely limits the transmission bandwidth of step-index multimode fibers, as is well known. A simple ray optics model of a ray bouncing off the core–cladding boundary at exactly θ_C gives an estimate of it,

$$\Delta v_m = \frac{c}{n_1} \left[1 - \frac{1}{\sqrt{1 + (\text{NA})^2/n_1^2}} \right] \approx \frac{c(\text{NA})^2}{2n_1^3}, \quad (8.19)$$

which for an NA of 0.2 predicts a velocity spread of over 1%—around 45 ns/km.

The exact modulation behavior of a given fiber depends on its length, manufacturing variations, and the details of its mode spectrum. We can feel pretty safe as long as the phase spread between modes is below a radian or so at the highest modulation frequency, but we expect to be in big trouble when it gets to half a cycle. That leads to a bandwidth–distance trade-off,

$$L \times \text{BW} < \frac{cn_1}{\pi(\text{NA})^2}, \quad (8.20)$$

which is pretty sharp; that 45 ns/km limits us to 3.5 MHz in a 1 km length. Bending the fiber into a coil reduces this effect and in rectangular cross-section waveguides can almost eliminate it; bending redistributes light between low- and high-angle modes, so that most of the light will spend some time in each, reducing the transit-time spread. (This is rather like the quenched mode coupling of Section 8.12.1.)

Gradient-index fiber works like a big long GRIN lens, ideally making rays at all angles see the same propagation delay. Waveguide effects and manufacturing errors limit

[†]This is of course a different use of the word *dispersion* than elsewhere in optics, where it refers to variation of n with λ .

its intermodal dispersion to 0.5–2 ns/km, which is much better, but still pretty big. Single-mode fibers don't have this problem, which (along with much lower loss and lower fiber cost) is why they are popular.

A single-mode fiber suffers *waveguide dispersion*[†] as well, because the fiber structure influences κ , so that k_z is a nonlinear function of k , and v_p slows down at long wavelengths. However, the quartz also has normal material dispersion, with n increasing at short wavelengths; for ordinary silica fiber, the two cancel out in the neighborhood of 1.3 μm , the zero-dispersion wavelength. You can send a fast signal a long way at 1.3 μm .

The loss minimum is at 1.55 μm , not quite the same as the dispersion minimum. In order to save money on repeaters, people have designed *dispersion-shifted fiber*, whose index profile is modified to reduce the waveguide dispersion enough to shift the zero-dispersion wavelength to 1.55 μm . There is also *dispersion-flattened fiber*, which is a compromise between the two and has low but not zero dispersion between 1.3 and 1.55 μm .

8.4 FIBER TYPES

8.4.1 Single-Mode Optical Fibers

Single-mode fiber is pretty well behaved stuff on balance; it will carry light long distances without doing serious violence to it, except for messing up its polarization, and it can be coupled into and out of good quality laser beams without great difficulty (once you get used to it). Once you get the light in, you can pipe it around corners, split it, combine it, modulate it, attenuate it, and detect it much the way you would with RF signals in coax. Due to the fiber's polarization degeneracy, the polarization of the transmitted beam is unstable, and because of its perfect spatial coherence, there are a lot of etalon fringes that have to be controlled with angled fiber ends, which is very awkward. Most of the single-mode fiber sensors we'll encounter later are limited by these twin problems.

The étendue of a single-mode fiber is the smallest of any optical component's—it's about the same as the $A\Omega$ product of a perfectly focused Gaussian beam, which is $\lambda^2/2$, on the order of 10^{-8} $\text{cm}^2\cdot\text{sr}$ at 1.5 μm . A laser is the only source that can cope with that.

8.4.2 Multimode Optical Fibers

Multimode fibers have a large V , which means that many modes can propagate; the number goes as V^2 , and is on the order of 200 for a step-index fiber. For the same value of V , a graded-index fiber will have half the number of modes of a step-index fiber. (Why?) These modes couple into one another at the slightest provocation, and as we saw, in step-index fibers they have very different propagation velocities, leading to severe bandwidth degradation.

Graded-index multimode fibers are the most common type; making the index of the core decrease smoothly from a broad peak on the axis brings the propagation velocities of the modes very close to one another, and so reduces the intermodal dispersion to a

[†]The term “mode dispersion” should be avoided, because different writers use it to mean both intermodal dispersion and waveguide dispersion.

few nanoseconds per kilometer. With laser light, the output of the fiber is a speckly mess that wiggles around like sunlight filtering through the trees on a windy day. It's even worse if the whole mode volume hasn't been excited, as often happens when we couple a laser into it with a low-NA lens.

On the other hand, multimode fiber is pretty stable with broadband illumination, if we carefully excite all the modes, for example, by using a diffuse source such as an integrating sphere, or homogenizing them by mandrel wrapping.

Multimode fiber is inherently lossier than single mode. High angle modes are weakly guided and so are vulnerable to scatter from irregularities in the fiber, or to bending and pressure. You see this as a gradual decrease of the measured NA of the fiber as its length increases. This hits a steady state eventually, since low angle light gets coupled into high angle modes, eventually spreading the loss evenly throughout the modes. The high loss encountered by high angle modes makes the far field pattern of a long piece of multimode fiber approximately Gaussian.

We often want to couple thermal light into multimode fibers, and so need to know something about their étendue. It's pretty easy to calculate for a step-index multimode fiber, because the acceptance angle is more or less constant across the facet, so $n^2 A \Omega' \approx (\pi a^2)(\pi(\text{NA})^2)$. Calculating the étendue of graded-index fiber is not so easy, because high angle rays near the edge of the core aren't well guided, but go into leaky modes. We can intuitively see that the étendue will be significantly less, because high angle rays can only enter near the fiber axis. A pure ray optics approach could calculate the NA as a function of radius and integrate up to $r = a$, but we'll content ourselves with arm waving here; it turns out that GI fiber has half the modes of SI.

At any given wavelength, each mode is fully coherent, and so we can imagine making a couple of multiple-exposure holograms that would transform each individual mode into a different focused spot. Thus each mode should contribute about the same étendue as a focused spot, that is, about $\lambda^2/2$. We thus expect the étendue of fiber with a given NA to go as

$$n^2 A \Omega' \approx \frac{\lambda^2 N}{2}. \quad (8.21)$$

The wisdom here is that to pipe low spatial coherence light around, step-index multimode is really the way to go unless you need the fast temporal response.

8.4.3 Few-Mode Fiber

The worst case is fiber with only a few modes, for example, telecom fiber used in the visible. There aren't enough modes for their spatial pattern to average out, and their coupling is a really serious source of noise because their propagation speeds are very different; thus the relative phase delay grows rapidly with z , leading to FM-AM conversion. Because the two modes will not in general have the same polarization, phase drift leads to severe polarization instability. Furthermore, the pointing instability of few-mode fiber is appallingly bad, with the beam wandering around by a degree or more depending on the phase delay.

8.4.4 Polarization-Maintaining (PM) Fiber

Polarization-maintaining (PM) fiber is SMF intentionally made asymmetric in profile to break the polarization degeneracy of circular-core fiber, and is seriously misnamed.

A built-in stress birefringence or (less often nowadays) an elliptical core makes the two polarizations travel at slightly different speeds. The speed difference is fairly slight; the beat length is a few centimeters in ordinary PM fiber to 1–2 mm in high birefringence (HiBi) fiber. As we’ve just seen, a short beat length means that the coupling of the two modes is incomplete. If we launch light into the fiber in one polarization state and look at what comes out, we find that the polarization is reasonably stable; regardless of mild abuse, bending, and so on, the polarization ratio at the output is typically 20 dB. That’s enough to help, but far from enough to fix the polarization noise and drift problem; a 20 dB polarization ratio means that the polarization axis can wander $\pm 6^\circ$, and you aren’t going to do a dark field measurement with that.

Furthermore, the short beat length means that the delay discriminator action is working moderately hard at demodulating the FM noise of the laser; a PM fiber has lots of noise if the orthogonal output polarizations are made to interfere, for example, by a misaligned polarizer or an angled fiber end, and the mode coupling eventually does the same thing to us.

The real reason that PM fiber is misnamed is that it doesn’t preserve any old input polarization, just the two eigenmodes. Its birefringence drifts with temperature, and the drift rate is proportional to the total retardation, so high birefringence PM fiber (1 mm beat length) with equal amounts of light in both modes has around *three orders of magnitude more* polarization drift with temperature than ordinary single-mode fiber.

We therefore use PM fiber with as pure a polarization mode as we can give it. Since the output polarization purity is typically 20 dB, using PM fiber that way makes the small-scale instability worse, but restricts it to a limited range of elliptical polarizations about the desired mode.

The great virtue of PM fiber is thus to stabilize the operating point of our polarization-sensitive sensors well enough that a properly normalized AC measurement can pull good data out without needing active stabilization, and that’s truly a major point gained—we don’t have to ship a graduate student with each instrument.

8.4.5 Chalcogenide Fibers for IR Power Transmission

A silica fiber will transmit at least some light from 250 to 2000 nm, and fluoride fibers reach 3000 nm, but the loss gets pretty large near the extremes. In the mid-IR, good fibers can be made from chalcogenide glasses. The main types are sulfide glasses, $\text{As}_{40}\text{S}_x\text{Se}_{60-x}$, which work from 1 to 6 μm , and telluride glasses, $\text{Ge}_a\text{As}_b\text{Se}_c\text{Te}_d$, for 3 to 11 μm . These are flexible enough, and low enough in loss, to be an attractive way of piping mid-IR power around, for example, CO_2 light for heating or surgery. Even at their best, chalcogenide fibers’ losses are between 0.2 and 1 dB per meter, thousands of times worse than single-mode silica at 1.5 μm , but they can handle CW power levels of 100 kW/cm² and pulsed levels of 200–400 MW/cm².

8.4.6 Fiber Bundles

The relatively low cost of fiber lends itself to some creative brute force solutions. For example, say we want to be able to pipe light from a lamp and condenser around easily. The étendue of a multimode fiber may be 10^{-5} cm²·sr, but we can use 10^3 of them in

a bundle, which starts to be quite respectable (and no longer so cheap). Bundles work best with broadband illumination, since otherwise speckles due to mode shifting from bending and stress will drive you crazy. They come in all sizes from two fibers to 20 mm diameter or larger; it depends on how much you want to pay for material, and for all that laying, cleaving, and polishing. They are always multimode because of the ratio of the core and cladding areas, but cladding modes are important and may need to be removed by potting in black wax.

The basic distinction is between image bundles, where each point on the output surface connects to the corresponding point on the input, so that an image can be transmitted, and random bundles, where no such care has been taken. Random bundles aren't really very random; they won't homogenize a space-varying illumination function, for example. The fibers just go any old way they happened to land. Now that LEDs are so good, random-bundle illumination is much less useful than it once was—we might as well mount the LEDs at the business end and run wires instead.

Imaging bundles are further divided into fused and flexible types. A flexible bundle has the fibers glued together only at the ends, leaving it free to bend. Fused bundles are used mainly for manipulating images; a short bundle with one concave and one flat surface makes a fiber-optic field flattener (also called a fiber-optic faceplate), which gives image tubes a flat image plane. Two image tubes can be cascaded with a biconcave fiber plate, which works better than any lens system and is dramatically smaller too. Tapered bundles, made by softening the bundle and stretching one end, magnify the image by the ratio of the face diameters. You can even twist the softened bundle, to make a fiber image rotator.

Of course, the phase gets completely screwed up in the process of transmission—the fibers are multimode, and no two are identical anyway. None of these image bundles can produce an aerial image—if the image is out of focus on the entrance face, it's out of focus forever, just like a ground glass screen. That's why the common term *coherent bundle* is a misnomer. On the other hand, the bulk optics alternatives are either a long succession of relay lenses or a long, specially designed GRIN rod. These are no picnic to get right, on account of the buildup of field curvature in the lenses and the specialized nature of the GRIN rod, but most of all because they can't bend without fancy and delicate articulated mirror mounts at each hinge.

8.4.7 Split Bundles

You can get split bundles that allow you to couple both a source and detector to a remote head. This can be pretty useful sometimes; with a tungsten bulb on one end and a spectrometer on the other, it allows decent low resolution spectroscopic measurements in tight spaces, for example. Combining a split bundle with a light bulb, band-limiting filters, and a monochromator is easy and very successful. If the monochromator is actually an integrated grating spectrometer/OMA design that plugs into a PC, you can put a nice film thickness measurement device together inexpensively that will do measurements under water, in awkward places, and in the face of lots of vibration and shock.

Another application of bundles is in fiber optic slip rings. A modulated laser on a rotating part is focused onto a surface having an annular ring of fibers. All the fibers are gathered together into one photodiode at the other end. The focused spot radius is about one fiber diameter, so that communications don't drop out periodically as the spot rotates. Communication the other way can be done with a more powerful laser

feeding the entire bundle, with the photodiode on the rotating part, but this is seldom done because of the laser power required.

8.4.8 Coupling into Bundles

We usually use a condenser to launch light into a bundle, as we saw in Example 2.1. The output end is a bit dicier. Well-homogenized multimode fibers illuminated with white light have nearly uniform output brightness across the core, but it narrows and dims for long lengths and short wavelengths, due to the excess loss in the high angle modes.

The fibers themselves do not pack all that tightly (see the problems), so that the illumination function is inherently nonuniform. The cleaved fiber ends have some nonzero wedge angle, which is usually nonuniform among the fibers in a bundle, so the far field pattern is a bunch of more-or-less overlapping cones rather than a single cone. This can lead to artifacts in grating spectrometers, for example.

Many-fiber bundles are more predictable than few-fiber ones, because the holes in the illumination pattern are smaller by comparison with the whole bundle. A fiber bundle used to receive light from the apparatus subjects it to a very odd spatial filtering function, which is usually poorly understood, as shown in Figure 8.3. Don't do it if you don't have to, and if you do, think about putting the bundle far enough out of focus that the pattern doesn't cause artifacts. A 125 μm fiber has to be surprisingly far out of focus for the modulation of the intensity to be below 1%, say.

Fiber bundles are more or less telecentric, since the angular pattern from each fiber is more or less the same. You can get fiber ring lights, which go around the outside of a lens, and provide a nice uniform illumination with an odd annular pupil function.

8.4.9 Liquid Light Guides

A variation on the fiber bundle is the liquid light guide, which is just a huge flexible waveguide made from high index liquid in a low index plastic tube. They're flexible but have poor illumination patterns. These have been pretty well superseded by white LEDs mounted where the light is needed.




Figure 8.3. Pupil function of a fiber bundle.

8.5 OTHER FIBER PROPERTIES

8.5.1 Leaky Modes

Modes that are just allowed or just cut off are generically called *leaky modes*. They get excited any time there's a mode mismatch. Their losses are very high, of course, which makes the overall fiber loss look anomalously high for the first bit of fiber after a launcher, splice, or connector.

8.5.2 Cladding Modes

Even leakier are *cladding modes*, in which the jacket takes over the role of the cladding and the whole fiber that of the core. You'll get these nearly always, but they're especially troublesome with single-mode fiber because the étendue of the highly multimode-fiber/jacket system is far larger than that of the core/cladding system. Cladding modes can be eliminated in a short distance by stripping the fiber jacket and putting a blob of index-matched black wax on the cladding, or by mandrel wrapping.

8.5.3 Bending

Fiber guiding occurs because there is no propagating mode in the cladding whose k_z phase matches with the light in the core. As soon as we bend the fiber into an arc, though, that is no longer true; light going round the arc at a larger radius sees a longer path length, so it can move faster and still stay phase matched. For any finite bend radius, there is a distance into the (notionally infinite) cladding at which a propagating mode is phase matched.

For long bend radii, that isn't a problem, because the guided mode has died off more or less completely before reaching that distance. As the bend gets tighter and tighter, though, the phase matched distance gets closer and closer, until it starts overlapping the propagating mode significantly and light starts getting lost. At that point, the loss increases exponentially as the bend radius declines. You can easily see it by putting a HeNe beam through a fiber and bending it with your thumb and forefinger.

For a given fiber, the guiding gets weaker as λ increases. The *bend edge* is the longest wavelength for which guiding occurs, as a function of the radius of the bend. You normally won't get very close to it.

8.5.4 Bending and Mandrel Wrapping

Bending loss has its uses, especially in stripping unwanted cladding modes and filling up the mode volume in a multimode fiber. It's pretty safe to bend a fiber into a coil 600 fiber diameters across, and you can go as small as $200\times$ briefly. For lab use, where broken fibers are easily replaced, you can be much more daring; a $125\text{ }\mu\text{m}$ OD fiber can be wrapped 10 times around a pencil to strip leaky modes and cladding modes, for example. In an unattended system, stick to the $600\times$ rule and use lots more turns, or use the black wax trick for cladding mode stripping instead.

8.5.5 Bend Birefringence and Polarization Compensators

Wrapping a fiber around a mandrel makes it birefringent, with the fast axis normal to the axis of symmetry. How birefringent a given piece of fiber becomes at a given wavelength

is usually a pretty well kept secret—you have to measure it. A rule of thumb is that the retardation in meters of a single-mode fiber wrapped into one turn of diameter D is

$$\beta_b = K_\lambda \frac{d^2}{D^2}, \quad (8.22)$$

where $k \approx 0.13$ m at 633 nm for silica fiber; a 5/80 μ m HeNe fiber looped into one turn of 72 mm diameter makes a nice quarter-wave plate, and that in turn allows us to make all-fiber polarization compensators.

Just as two quarter-wave plates can turn any fully polarized state into any other, so two loops of fiber attached to flat discs, hinged so as to allow them to rotate through 180° , will allow us to control the polarization funnies of our fiber, at least until it starts moving again. Three-paddle compensators (two $\lambda/4$ and one $\lambda/2$) are often used to provide some tolerance for errors and wavelength shifts, because if your wave plates are not exactly $\lambda/4$, there are some places you can't get to—just the way you can touch your nose to your shoulder and to your wrist, but not to your elbow. Since we're usually interested in linear polarization, after tweaking the two $\lambda/4$ paddles to get rid of ellipticity, we can get any linear polarization by turning the $\lambda/2$ paddle.

8.5.6 Piezo-optical Effect and Pressure Birefringence

Squashing a fiber produces a piezo-optical shift in n , as we saw in Section 8.5.6. Quartz isn't very piezo-optic, but there's a lot of it in your average fiber system—and even a random effect grows with \sqrt{L} .

8.5.7 Twisting and Optical Activity

Analogously, if we twist a fiber about its own axis, it becomes optically active. The total polarization rotation of a fiber twisted through ξ radians is

$$\Delta\theta = g\xi, \quad (8.23)$$

where g is about 0.15 for silica fiber at 633 nm.

Aside: Normal Propagation. When we discuss the propagation of light in circular core, single-mode fibers under peculiar conditions (e.g., bending, twisting, and so on), we are looking at the deviation of the light from the normal situation, that is, \mathbf{k} and \mathbf{E} being constant in the lab frame. Twisting a fiber drags the polarization direction along with the twist, but that's solely a material property. Maxwell's equations and the waveguide properties *per se* aren't altered, so without the material change, the light would continue with the same polarization in the lab frame, completely oblivious to the change in the guide. Things are a bit more complicated with PM fiber, because we have to look at the twist or bend as a coupled-modes problem.

8.5.8 Fiber Loss Mechanisms

Silica fibers are amazingly low loss devices, far more transparent than air in some spectral regions. Loss is dominated by Rayleigh scattering and microbending at short wavelengths, although there is also strong electronic absorption deeper into the UV.

The Rayleigh scatter goes as λ^{-4} , so it's very strong below 400 nm and almost nonexistent beyond 1.2 μm . There are overtones of the OH vibrational resonance at 2.73 μm , which occur in bands near 950, 875, 825, and 725 nm, and molecular vibrational absorption beyond $\sim 2 \mu\text{m}$. The sweet spot is in the 1–1.6 μm range, where the absorption, Rayleigh scatter, and microbending are all weak, and the fiber can achieve losses of 0.15–0.5 dB/km. By comparison, a 1 m thick piece of window glass looks as green as a beer bottle (try looking through a bathroom mirror edgewise sometime). The few meters of silica fiber in the typical instrument will transmit light from 250 to 2000 nm or thereabouts.

At 254 nm, silica fiber's absorption goes up to around 1000 dB/km, and permanent darkening or *solarization* may start to occur due to UV damage. You can get solarization-resistant fiber that works stably down to 180 nm (1.5 dB/m). Multimode fiber is especially flaky in the UV, because the mode coupling is very strong. Low angle modes get coupled into high angle modes, which are attenuated much more strongly. Wiggling the fiber makes the amplitude go up and down, by as much as 1 dB/m at 254 nm.

Minor manufacturing errors also contribute to fiber loss; small-scale roughness at the core–cladding interface (a form of microbending) is the major one, but there is also the occasional bubble or inclusion. Microbending is worst at high Δ , just as high index lenses are more vulnerable to surface errors. These intrinsic losses dominate at long lengths, but in instruments we're more often fighting splices, connectors, and larger scale bending.

8.5.9 Mechanical Properties

Glass is a brittle material, which is to say that cracks that are not held shut will eventually propagate through it, causing failure. The time to failure goes as a very high negative power of the tensile stress. At the crack tip, chemical bonds are being broken, which takes energy. It takes a lot less energy if there's water around to bond to the newly created surface, so the time to failure is enormously longer for dry fiber versus wet.

It's odd at first sight, but the strength of a fiber depends on its length. For a uniformly loaded fiber, it's the largest flaw in the entire length that causes failure, just like the weak link in a chain. The statistics of how big the largest flaw is are sensitive to how long it is, and of course very sensitive to the history of that piece of fiber.

The fracture strength of fibers thus varies all over the place, with some published curves showing 0.1% probability of fracture per kilometer at stresses of only 50 kPa, which is less than the stress on your spine from holding up your head. On the other hand, a really good piece of fiber breaks at around 4 GPa, which is three times stronger than the best steel.

8.5.10 Fabry–Perot Effects

In Section 11.9.2, we'll see that minor misalignment protects us from etalon fringes to a considerable degree. There's no way to misalign a single-mode fiber, so we would seem to be in trouble; along with the polarization instability, etalon fringes are the leading cause of trouble in laser-based fiber instruments. One thing that helps is to cut the ends of the fibers at an angle, so that the reflection is outside the fiber NA, and so goes off into the cladding to be absorbed. Fibers tend to cleave normal to the axis, but twisting them while cleaving will produce an angled facet. Angled cleavers exist that work reasonably

repeatedly, and you can get angled physical contact (APC) connectors and matching fiber polishing kits. These will get you down to around 10^{-4} reflectance, which is still poor in the scheme of things, but is pretty good for fiber. We're often forced to choose fragile APC connectors and expensive Faraday isolators, which significantly reduces the attraction of fiber.

One gotcha is that the etalon fringes in fiber are polarization dependent, because different polarizations see different delays. In a broadband system, this will decorrelate the noise in the two polarization eigenstates.

8.5.11 Strain Effects

If you imagine stretching a hollow metal waveguide to $1 + \epsilon$ times its original length, it's easy to see that the two main effects are a propagation distance increased by the same factor, and a decreased propagation speed due to the slight narrowing of the guide due to the stretch.[†]

Fiber behaves the same way, with the addition of the change in n due to strain. The rate of change of phase with respect to stretching in a fiber is

$$\frac{d\phi}{dL} = nk_0\xi, \quad (8.24)$$

where ξ is given by

$$\xi = 1 - \frac{n^2}{2}(P_{12} - \mu(P_{11} + P_{12})) \approx 0.78 \quad (8.25)$$

so that

$$\frac{d\phi}{dL} \approx 0.78nk_0. \quad (8.26)$$

8.5.12 Temperature Coefficients

We saw in Section 4.2.2 that the temperature component of the optical path length is


$$\frac{\text{TC}_{\text{OPL}}}{\text{OPL}} = \frac{1}{n} \frac{\partial n}{\partial T} + \text{CTE}. \quad (8.27)$$

The CTE of fused quartz is very small, about $5 \times 10^{-7}/^\circ\text{C}$, but its TCN is around $+9 \times 10^{-6}/^\circ\text{C}$, so that its TC_{OPL} is $7 \times 10^{-6}/^\circ\text{C}$.

Although the exact value is slightly modified by waveguide effects, this is the dominant effect on the phase change in fiber devices with temperature, and also in the temperature tuning of fiber Bragg gratings.

If the fiber has a strong plastic cladding (e.g., 1 mm diameter nylon), the temperature sensitivity will go up by an order of magnitude due to the huge CTE of the plastic straining the fiber.

[†]Stretching things makes their dimensions perpendicular to the strain shrink. Poisson's ratio μ tells how big the relative shrinkage is. For something like rubber or chewing gum, where the total volume stays constant, $(1 + \epsilon)(1 - \mu\epsilon)^2 = 1 + O(\epsilon^2)$, so $\mu = 0.5$. Glass has $\mu \approx 0.4$, and most metals, 0.33.



$$E = \sum_{n=0} t_1 t_2 (r_1 r_2)^n e^{i 2n \Delta\phi}$$

Figure 8.4. Demodulation of laser noise by etalon fringes and double Rayleigh scatter: this is generally what limits the SNR of fiber measurements.

8.5.13 Bad Company: Fibers and Laser Noise

We saw in Section 2.5.3 that multiple reflections can cause serious intensity noise. As shown in Figure 8.4, even a short length of fiber can do it, and so can multiple Rayleigh scattering even in a perfectly terminated fiber. This coherence effect is one of the main SNR limitations of fiber sensors.[†]

If the scatterers are localized (e.g., facet reflections), you get fringes in the spectrum from localized mirrors. Even tiny frequency fluctuations get turned into huge amplitude variations due to the very steep slope of the fringes with frequency.

In the double Rayleigh case, the phase of all the little scatterers is random, so the scattering contributions add in power. Either way, for fibers longer than $1/(n\Delta\nu)$, you basically get the whole interference term of the scattered light turning into noise with a frequency spectrum that is the autocorrelation of the source spectrum. There's the usual 3 dB reduction in intensity noise, as in Section 13.6.9, since some of the interference term remains phase noise instead of all of it becoming amplitude noise. As we saw in Section 2.5.3, this effect can become dominant with surprisingly short path differences.

The fringe phase is strongly dependent on fiber length and stress birefringence, so two fibers fed from the same source will produce a lot of uncorrelated noise. This effect makes double-beam noise reduction systems such as laser noise cancelers (see Section 10.8.6) far less effective with fibers.

8.6 WORKING WITH FIBERS

8.6.1 Getting Light In and Out

Getting light out of a fiber is relatively simple, because the light is already coming out on its own; you just stick the facet at the back focus of a collimating lens, and you're done. The NA of the fiber makes $10\times$ to $40\times$ microscope objectives a good choice; if you want good beam quality, use single-mode fiber, a good lens, and a good solid

[†]Amnon Yariv et al., Signal to noise considerations in fiber links with periodic or distributed optical amplification. *Opt. Lett.* **15**, 1064 (1990).




Figure 8.5. Coupling to fibers.

mount—not some big six-axis stage. With multimode fiber, there’s no avoiding an ugly beam, so don’t worry about the lens quality so much.

If you just want to dump the light, or shove it into a photodiode, you can dunk the fiber end into some index matching material (oil, gel, or wax). If you get the index right, the back reflection will be very small, and virtually none of the escaped light will bounce back into the fiber—the low étendue works in our favor here.

Getting light into the fiber in the first place is much more difficult, because you don’t have the light itself to guide you (see Figure 8.5.). The first thing you need is to believe in reciprocity. Reciprocity says that an antenna used for receiving has the same spatial pattern and the same gain as when it’s transmitting, or that a beam coupling into a fiber has the same loss that you’d see if it were coming out of the fiber and being spatial-filtered to look just like the incoming beam.

Thus it is necessary to mode-match to get the best coupling efficiency. Because the beam coming out of the fiber is nearly Gaussian, with an NA specified by the manufacturer, you can just use the same NA on the launching side, and be sure of the best coupling.

If your coupling efficiency is below 80%, you haven’t got it quite right; measure the NA of the output of the fiber (remember we’re talking about the $1/e^2$ diameter), and use that on the input side. Your coupling efficiency should improve to the 80–90% range. The coupling efficiency as a function of mode field diameter mismatch is roughly

$$\eta_c \approx \frac{2w_1w_2}{w_1^2 + w_2^2}. \quad (8.28)$$

Focus error and aberrations do to coupling efficiency just what they’d do to an interferometer, and for the same reasons. There are formulas available to calculate how close they should be, but basically you need to be within, say, 0.1 Rayleigh ranges of perfect focus over temperature. This isn’t particularly onerous, since the Rayleigh range is tens of microns.

The major aberration we have to deal with is astigmatism in diode lasers, and even this isn’t normally such a big problem since the diode’s beam is oblong and we’re only using the central bit, where the aberration hasn’t had time to get too large; the coupling efficiency is dominated by the shape mismatch unless you circularize the beam. There are formulas for calculating all these things, but nobody ever uses them in real life.

Aside: A Slightly More Rigorous Look at Reciprocity. An electromagnetic definition of a reciprocal component is that its behavior is unchanged when we replace t with $-t$, \mathbf{E} with \mathbf{E}^* , and \mathbf{H} with $-\mathbf{H}^*$. Consider shining a laser beam through a component, and putting a lossless phase-conjugate mirror after the component, so as to send the beam back through it the other way. The component is reciprocal if and only if the two-pass beam is a phase-conjugated replica of the incoming beam, independent of polarization. Lossless lenses, mirrors, wave plates, beamsplitters, and fibers are reciprocal; attenuators, polarizers, and Faraday rotators are not. Loss can be put in by hand, so we often loosely think of attenuators and polarizers as sort-of reciprocal, though a specific term for that would be useful. Hand-waving reciprocity is a very useful working concept.

8.6.2 Launching into Fibers in the Lab

Those nice fiber collimators that Thor Labs and others sell are great for coming out of fiber, but much less good for going in, because you can't see what you're doing during alignment very well. If you don't mind searching blindly for some time, you can put one of these in a two-axis tilt mount and twiddle till you see something. Alternatively, if yours are connectorized, you can start with a multimode fiber patch cord running into an optical power meter. Find the points where the intensity falls off, and then adjust so you're halfway in between. Then when you put your single-mode patch cord on, you'll see some light right away. For visible light, you don't need the power meter—provided the total power is low enough to be eye-safe, just look at the end of the fiber.

Sometimes nobody makes collimators suitable for our wavelength or beam diameter. The wavelength problem is due to chromatic aberration in the coupler lens, which we can fix, and mistuned coatings, which we can't. The power of lenses generally increases toward shorter wavelength, so if your coupler's center wavelength is too short (e.g., using a 1310 nm collimator at 1500 nm), try a weak convex lens in front of it, and if it's too long, try a weak concave one.

For the fully manual approach, the first thing to start with is a 20 \times , 0.4 NA microscope objective. That's far higher than the NA of most fibers, but then most laser beams are far smaller than the pupil of the microscope lens, so the actual NA you wind up with is closer to 0.2 in most cases anyway. Using a 10 \times , 0.2 NA will require filling the whole pupil—5 mm in diameter or more—and doesn't leave any elbow room for the wings of the Gaussian. Here are the steps (refer to Figure 8.6).

1. If the laser is visible, measure its power and adjust it with an ND filter until it's well below the Class II limit (200 μW is a good number); the most sensitive way to do the early aiming is by looking into the fiber, and it's nice to not hurt ourselves in the process. (This applies in fault conditions too; for example, you bump the power control knob or the monitor photodiode lead comes loose. If you don't have a really trustworthy power meter, don't look into the fiber.)
2. Get the laser beam running horizontally, parallel to one edge of the optical table, and at a convenient height (about 150 mm). Mark the wall where it hits.
3. Strip and cleave the fiber to get a clean end with 1 or 2 cm of stripped length, and have a good look at it with a fiber microscope. Put it in a holder, either a ceramic V-groove collet or a milled inside corner, and clamp or wax it in place.
4. Put the fiber on a three-axis mount with locking screws, and check with a ruler that the axis of the fiber chuck is horizontal and that both it and the z translator




Figure 8.6. Launching light into a fiber.

are running parallel to the table edge ($\pm 1^\circ$ or so is OK). Check that the light beam hits the fiber end.

5. Put a 20 \times microscope objective on a sturdy mount, and move it until the center of the output beam coincides with the spot on the wall. Look at the shadow of the fiber in the light beam when it's backed off—the shadow should stay centered and not walk off sideways as z is adjusted.
6. Put an optical flat (or even a microscope slide, in a pinch—AR coated, preferably) on an ordinary two-tilt mirror mount between the laser and the objective. (This is a vernier adjustment, because the amount of tilt you can get isn't too big, and since it only tilts the output beam and doesn't translate it, it should be orthogonal to the xyz adjustments.)
7. Watch the way the light reflects from the fiber facet as you adjust z to get it centered, a millimeter or so past focus. Light should be visible coming from the fiber.
8. Twiddle the xyz knobs to maximize the light coming out of the fiber end. It probably won't be too impressive yet because we haven't adjusted the angle carefully.
9. Adjust the optical flat tilts to maximize the transmitted power. Iterate the xyz and angular adjustments until the coupling is maximized.

If you're using ordinary SM fiber, you're done. It's a bit tougher with PM fiber, because we have to keep the polarization pure going in, and we don't in general know the orientation of either end of the fiber. For PM fiber, add these steps.

PM Fiber. Before adjusting,

- 4b. Make sure the polarization going in is really linear.
- 4c. Put a rotatable analyzer at the output of the fiber.

- 5b. Use a zero-order quartz $\lambda/2$ plate instead of the flat; it's thick enough to do both jobs if we're careful. You can use both if you like, or if your wave plate is too skinny to displace the beam much without a huge tilt.
9. Iterate rotating the analyzer to minimize the transmitted light and rotating the wave plate to minimize it some more.
10. Jiggle the fiber and look at how much the intensity jiggles. Iterate some more until the jiggle is as small as it's going to get.
11. If you're using a diode laser, you can use minimum detected noise on a photodiode as your criterion for polarization twiddling instead of minimum transmission.
12. Mark the axis orientation of the fiber at both ends for next time.

8.6.3 Waveguide-to-Waveguide Coupling

The easy way to couple two fibers is to use a connector splice. To avoid the resulting broad etalon fringes, you can use microspheres, little spherical lenses that you can put between two fibers. This isn't too easy to do manually. GRIN rods are perhaps a bit more useful; a GRIN rod of $\frac{1}{4}$ period or a bit more will couple two fibers together very nicely, and you can increase the NA by going closer to half a period, at the expense of working distance. Stick the GRIN rod in place, position the first fiber so that light comes out with the same NA that it went in (easy to see on a white card a little way away), and then go through the steps in the previous section.

Connecting fibers to diode lasers and integrated optics devices is usually done in one of four ways: microspheres, lensed (i.e., domed) fiber ends, GRIN rods, or proximity (butting the fiber right up to the device). GRIN rods and proximity are the two easiest ways if you're rolling your own.

8.6.4 Connecting Single-Mode to Multimode Fiber

You can easily send light from a single-mode fiber into a multimode one, but not the other way. In general, trying to get light from multimode into single mode will cost you about 20 dB in signal, together with very large (order-1) variations due to bending and temperature changes. This is because your average multimode fiber has about 50–100 modes, all of which are orthogonal. Because of their orthogonality, they can't all couple well into your fiber mode—in fact, if all N modes are illuminated equally with incoherent phases, you can get at most $1/N$ of the light into your single-mode fiber.

8.6.5 Fibers and Pulses

One situation where fiber really does behave a lot like coaxial cable is when you're using short pulses with a very low repetition rate. A single pulse entering a complicated fiber system will generate all sorts of smaller pulses due to reflections, which will rattle round the system for some time. On the other hand, if the pulse is short compared with the delays between different reflections, and the rep rate is low enough for them all to die away before the next pulse, things can be pretty well behaved. The author has a 20 Hz, 20 ps tunable laser system, which works fine with fibers as long as the reflections are suppressed by time gating. Of course, the polarization instability problems remain.

8.6.6 Mounting Fibers in Instruments

In instruments, we have another set of problems, caused by dirt, vibration, shock, and temperature drift, and mechanical instabilities due to lubricant flow and stick–slip energy storage (you know, like the San Andreas Fault). It isn't a problem inside the fiber, obviously, but it is for launching and collimating. The easiest and most expensive way round this is to start with pigtailed components and just stick them together with connectors. You can get armored fiber patch cords, whose steel gooseneck jackets enforce a minimum bend radius and can be bolted down, which helps a lot. The corrugated interior causes microbending, though, so they aren't as useful with multimode fiber, especially in the UV.

One thing to remember is that connectors for weird fiber may be impossible to find (and the price may wish you hadn't found them if you do), so try to stick with 125 μm cladding diameters which is the standard for communications fiber.

The other thing about using fibers in instruments is that you have to be very careful in choosing a jacket material. Teflon and polyimide are the most common choices at present, and they're both good as long as they're not thick enough to stretch the fiber. We've already talked about a thick plastic jacket straining the fiber with temperature, but there are other possible problems. Nylon and some other plastics are hygroscopic, so that they swell in humid conditions—you don't want that source of instability either.

8.6.7 Connectors

Fiber connectors used to be expensive and very miserable, but now they're cheaper and more mildly miserable. (At least if your fiber has a 125 μm cladding OD—special sizes are much more expensive, (for example, \$40 for an 80 μm ST connector, versus \$8 for 125 μm .) This is of course a jaundiced view, but how else do you describe a splice that is guaranteed to send a few tenths of a percent of the light right back into the source (or 0.01% for angled ones)? Even that 100 ppm is more than enough to make a diode laser misbehave and will cause percent-level etalon fringes.

8.6.8 Splices

The lowest-reflection way to connect two fibers is fusion splicing, where two cleaved fibers are butted and then melted together, but that isn't too convenient. We more often use epoxy splices, which come in kits and are easy to use. Fusion splices have reflections of 1 ppm or so, which is a big help in reducing instability. There are also temporary splices, mostly based on very accurate V-grooves in ceramic materials, with a variety of clamping strategies. You stick the fibers in, usually with a little index gel, and engage the clamp. Expect 0.1 dB coupling loss for a fusion splice and 0.3–0.5 dB for an epoxy splice.

Remember the egregious polarizing beamsplitter cube in Example 4.1, whose reflections were 1%, and whose length was only 25 mm—it had a temperature coefficient of transmission of 12%/°C. A fiber patch cord is even worse, at least with highly coherent sources, and you even get strong etalon fringes between the two fiber ends in the connector itself. As usual, you get strong fringes with single-mode fiber because there's no misalignment to help suppress them.

You'll probably use FC or ST connectors, in either normal or physical contact (FC/PC, ST/PC) varieties. The PC types use fibers with a slightly convex polish on the ends; they

have lower loss and back-reflections, but they are easily destroyed by dust on an end facet. You can also get angled physical contact (APC) connectors, with a nominal 8° facet angle, and these help quite a bit with back-reflections, as long as the reflected wave is all beyond the core/cladding critical angle, so that it all leaks out. APC connectors are unforgiving of angular misalignment, so they come with keys to keep them straight. These keys are the same kind used in PM fiber connectors, which can make life interesting if you're not careful.

You can't mate an APC connector to another kind, because the light comes out at a 4° angle to the fiber axis, and because you can't get the ends close enough together on account of the angle. That 4° angle makes trouble in coupling light in and out, too.

Connector manufacturers sell inexpensive kits that help a lot with getting the end facet polish right, so by all means buy one. You can clean fiber facets using a Q-tip and some very clean ethanol, or (for quick and dirty use) some scotch tape.

8.6.9 Expanded-Beam Connectors

At the price of a couple of decibels' loss per pair, you can use connectors with integral collimating lenses. They are tolerant of dirt and other sorts of abuse, and you can shoot the beam some distance in mid-air if you need to, which is a plus in instruments. The lenses aren't good enough to make a decent beam, unfortunately.

8.6.10 Cleaving Fibers

Cleaving fibers is easy but takes a bit of practice; most of us are much too rough at first. Semiautomatic fiber cleavers exist, but careful hand work is just as good. Strip the fiber back a couple of inches, wet it, and nick it ever so gently with a V-point diamond scribe in the middle of the stripped section (don't buy the cheap carbide scribes). Bend the fiber gently at the nick until it breaks.

Another technique is to strip it back a bit further, wet and nick it as before, and bend it back in a hairpin loop between thumb and forefinger. Pull gently on one end so that the loop slowly gets smaller until it breaks.

Bad cleaves are usually due to too big a nick or to pulling the loop in too fast; hackles[†] form when the fracture propagates faster than the speed of sound in the glass. Big fibers are much harder to cleave than little ones. Twisting the fiber makes an angled cut.

8.7 FIBER DEVICES

It's worth emphasizing the distinction between all-fiber devices, which use some intrinsic property of the fiber, and fiber-coupled devices, where the fiber is grafted onto a bulk optic or integrated optic device. Directional couplers, Bragg gratings, and polarization compensators are all-fiber, but Pockels cells, isolators, and detectors are merely fiber coupled. All-fiber devices often have a cost advantage, whereas fiber-coupled ones are more expensive than bulk. Bulk devices are covered elsewhere, so here we'll stick with the all-fiber ones.

[†]A hackle is a bump or dip where a bit of glass has remained attached to the wrong side of the cleave.

8.7.1 Fiber Couplers

By bringing the cores of two single-mode fibers close together, the mode fields overlap, leading to the coupled-mode behavior we saw in Section 8.3.3. Phase matching makes such *evanescent wave couplers* highly directional. If the beat length is longer than the interaction zone, nearly 100% coupling can be achieved, but we more often want 3 dB, which gives maximum fringe visibility in interferometers. Couplers are usually made by fusing two fibers together and stretching the fused region until its diameter is slightly less than that of a single fiber, as shown in Figure 8.7a, which shows a 2×2 tapered coupler. By conservation of étendue, the NA must increase as the core diameters drop, and the high angle light is no longer guided. It therefore goes off into cladding modes. The reverse process occurs as the core broadens out again, corralling most of the light into the cores again. Because most of the light is in cladding modes at some point, the outside of the cladding must be clean and untouched; thus tapered couplers are usually suspended in air inside a case.

All these devices are inherently $2N$ -port, so in combining two beams with a 3 dB coupler, you're going to lose half of each one out the other coupled port. This is easy to see thermodynamically, because otherwise you could feed 100 fibers from a low temperature source, combine them, and use the result to heat a higher temperature sink. The four-port character is not necessarily a disadvantage, but you do have to be aware of it.

Coupler figures of merit are the coupling ratio, directivity, and excess loss. For a four-port (2×2 , i.e., two fibers in, two fibers out) coupler, if we feed P_1 into port 1, the coupling ratio is $P_4/(P_4 + P_2)$, the directivity is P_3/P_1 , and the excess loss is $P_1 - (P_4 + P_2)$.

Provided the taper is sufficiently gentle, the reflections are small, and the excess loss is well below 1 dB; large-core multimode splitters can have 1.5 dB of excess loss, and some single-mode ones reach 0.1 dB. The directivity is usually -50 to -60 dB (optical), and back-reflections are relatively low too, around -50 dB, because the fiber stays more or less intact. If you need adjustability, you can imbed fibers in epoxy blocks and lap




Figure 8.7. Fiber couplers: (a) tapered fused coupler and (b) lapped coupler.

them until the mode fields reach the surface; putting two such fibers together makes an adjustable coupler.

8.7.2 Fiber Gratings

A single-mode fiber is an ideal application for Bragg gratings. For a given wavelength, we know exactly what k_z is, and it's sufficient to put in a grating with $\mathbf{k}_G = 2k_z\hat{\mathbf{z}}$ to make a very effective mirror; in wavelength terms, the Bragg wavelength λ_B is

$$\lambda_B = 2n\lambda_G. \quad (8.29)$$

As with plane gratings, the resolving power of a weak Bragg grating ($R < 10\%$) is $R = mN$, the number of lines in the grating times the diffraction order (providing, of course, that the grating is weak enough that the far end gets illuminated). A weak Bragg grating exhibits a strong, narrow spectral dip, which is insensitive to polarization (see Figure 8.8).

You can use fiber Bragg gratings like mirrors, for example, to make fiber lasers or interferometers. You can use them as filters, say, to confine the pump light in a fiber amplifier or laser. And you can use them as fiber sensors, because stretching the fiber or changing its index will cause the reflection peak to move.

In the beginning, people made these fiber Bragg gratings by etching the cladding, but nowadays they're made by the photorefractive effect: strong UV (248 nm) irradiation causes

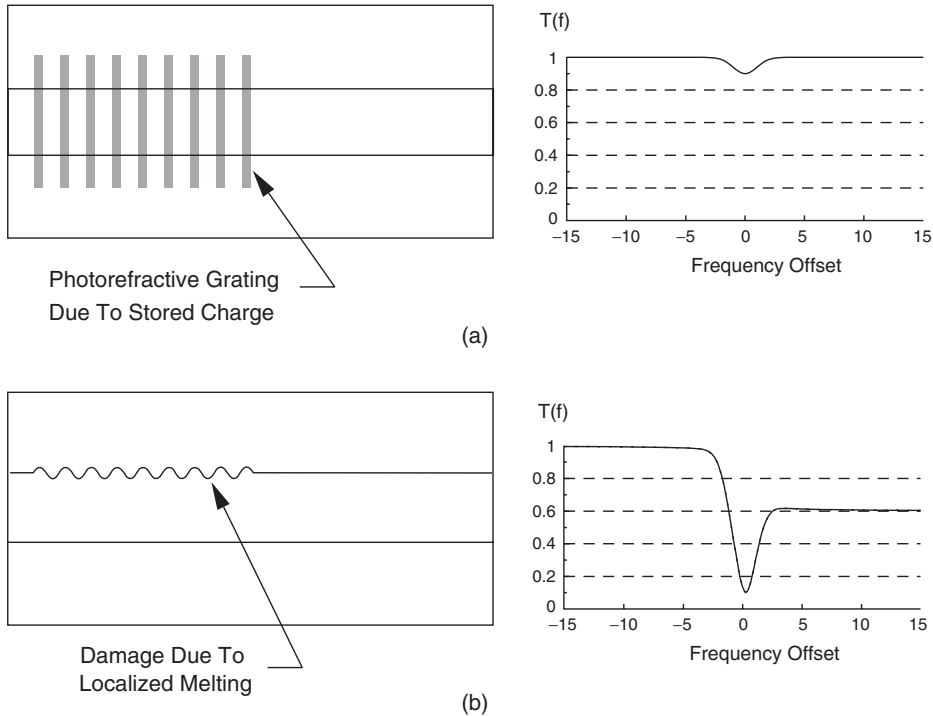


Figure 8.8. Fiber Bragg grating: structure and spectral behavior in reflection and transmission.

permanent changes in the index of the glass. Ordinary *Type I* fiber gratings (Figure 8.8a) have a small index change that extends more or less uniformly through the core, which makes the coupling to cladding modes very weak, normally a very desirable property.

A short-pulse, high rep rate excimer laser shining through a mask writes gratings in a fiber fed past from reel to reel. This is potentially very inexpensive in large quantities, although if you buy one unit, it isn't much different from a 25 mm square plane grating at this writing (\$500).

The peak reflectance of a Type I fiber Bragg grating is typically below 10%, though the newer commercial ones are long enough to reach 90% or more. The tuning of the grating depends on strain and temperature; the interaction between the two is slight, so the shift of the peak can be predicted from (8.24) and (8.27).

8.7.3 Type II Gratings

It is also possible to make extremely strong Bragg gratings, in which the grating lines are conceptually more like a stack of dielectric films; these Type II gratings are made by actually melting the core-cladding boundary with much higher laser fluence. They are asymmetrical and hence couple light into cladding modes very strongly for wavelengths shorter than λ_B ; the loss can be as much as 90% (Figure 8.8b).

Like dielectric stacks, these gratings can achieve very high, flat reflectivities over a moderately broad band—99.8% has been reported.

8.7.4 Fiber Amplifiers

Fibers doped with rare-earth ions are widely used as amplifiers; erbium at 850, 990, and 1550 nm, neodymium at 1.06 and 1.32 μm , and holmium at 1.38 μm . They boost the signal level without needing to be converted to electrical signals and back again, as in a repeater. The spontaneous emission increases the noise, so you can't use too many in a row. The spontaneous emission has shot noise, and it also leads to severe *beat noise*, caused by its interference with the desired signal. This is similar to the coherence fluctuations and double Rayleigh scattering problems, and shows up because there's no spatial averaging to get rid of it as there is in bulk optics. You can't do much about signal/spontaneous emission noise, but a narrow filter will mostly eliminate spontaneous/spontaneous beats, which can be much worse.

Erbium-doped fiber amplifiers (EDFAs) are pumped with special diode lasers at 980 or 1490 nm; the pump light is fed in via a coupler and stopped with a Bragg grating, or a coated facet, or sometimes just allowed to exit along the output fiber, where it will ultimately be attenuated by the fiber loss.

Don't be tempted to see fiber amplifiers as equivalent to packaged RF amplifiers in coax systems, because the resulting systems are not nearly as stable; a garden-variety RF amplifier gives you 20 or 30 dB of isolation from output to input, whereas a fiber amplifier gives none whatever. In fact, it gives less than 0 dB isolation, since the gain is not intrinsically directional—it amplifies the light going the other way just the same amount. RF folks would say that a 20 dB EDFA has *minus* 20 dB of isolation, which would make them unhappy, and for good reason. Part of the skill of designing fancy fiber systems is knowing how to use the minimum number of expensive isolators to get good stability.

8.7.5 Fiber Lasers

By adding Bragg mirrors tuned to the signal wavelength, putting mirror coatings on the facets, or by turning it into a ring resonator with a coupler, a fiber amplifier can be made into a solid state fiber laser. The simplicity of the idea is its main advantage, but in principle it can be of great benefit in long path difference interferometers such as Fabry–Perot types, because the linewidth can be made narrower than the typical 10–100 MHz of a single frequency diode, and the coherence length correspondingly longer.

8.7.6 Fiber Polarizers

Polarizers are an example of something that's hard to do in fiber but easy in bulk. Persuading a fiber to actually absorb the unwanted polarization requires giving it anisotropic loss, for example, by lapping the cladding down to get to the mode field, and then putting on a metal coating to absorb the TM polarization. The TM mode excites surface plasmon waves in the metal, which are enormously lossy. Metal fiber polarizers can have open/shut ratios of 50 dB and losses of 0.5 dB.

You can get polarizing (PZ) fiber, in two types. One works by the single-mode equivalent of frustrated TIR; if the evanescent field falls off at slightly different rates for s and p polarizations in the cladding, then a localized cladding region with high absorption or low refractive index can cause one polarization to be attenuated differently from the other. The other kind works by the difference in the bend edge in the two polarizations; a few meters of carefully aligned coiled fiber sends most of the more weakly guided mode off into the cladding. Usually we just put a bulk polarizer before or after the fiber. Walkoff plates are especially popular for this because one beam is undeviated, which makes them easy to make and to align.

8.7.7 Modulators

We've already encountered integrated optic Pockels cells, which are the most common modulators used with fibers. Fiber phase modulators are often made by wrapping many turns around a large piezoelectric tube, but these *fiber stretchers* can't go much faster than 50 kHz. They also have a lot of bend birefringence unless you anneal the fiber in the coiled state (which is hard since the piezo won't stand the 800°C annealing temperature).

8.7.8 Switches

The most common type of intensity modulator or switch is the integrated optic Mach–Zehnder type, in which a phase modulator in one arm of a waveguide Mach–Zehnder interferometer causes a bright or dark fringe to occur at the output waveguide. Interferometers are always four-port devices, even though these ones don't look like it; the third and fourth ports are radiation into and out of the substrate when the two waves are out of phase. You can also do it by putting fibers on piezos, and moving them in and out of alignment. People have built N -way switches in that fashion, but it's a hard way to make a living.

8.7.9 Isolators

Fiber isolators are also extrinsic devices; the two most common are Faraday and acousto-optic, which we're familiar with from bulk optics. Faraday isolators rely on polarizers

to get their directional selectivity, so use PM fiber or mount the isolator right next to the laser, where the polarization is very stable. Faraday rotator mirrors provide some isolation as well, providing a polarizing beamsplitter is used to separate out the returned light.

8.8 DIODE LASERS AND FIBER OPTICS

Both of these technologies are wonderfully useful, but they require much more engineering than one at first expects. Neither is a quick fix for an implementation problem. Especially seductive is the ease of hacking up a connectorized fiber system: just twist the connectors together, and you're done.

The parallel between this and doing signal processing with a table covered in boxes and RG-58/U patch cords is real but not perfect; the main differences are that the length of the optical fiber is at least millions of wavelengths, whereas the coax is usually less than 1, and that electronic amplifiers provide isolation, whereas almost all optical components are inherently bidirectional. An optical fiber setup is more closely analogous to a whole cable TV or telecommunications cable system, with miles and miles of coax. Stray cable reflections can cause severe ghost images in TV reception, much the way etalon fringes foul up measurements.

Even there, the analogy breaks down somewhat; the reflections from the cable network don't cause instability or frequency pulling of the sources, because they're well isolated; diode lasers are much more susceptible. And besides, coax doesn't have polarization or mode problems, or lose signal when you bend it. There's no coax network long enough to get coherence fluctuations from a typical oscillator with a 1 Hz linewidth, either.

Diode lasers and fiber optics work better apart than together, unless you're prepared to spend some real dough (\$1000 or more) for decent isolators, or to destroy the coherence by large UHF modulation, and even then, you have the coherence fluctuation problem we've already discussed.

8.9 FIBER OPTIC SENSORS

It is important to distinguish a fiber sensor from a sensor with fibers in it. There are right and wrong reasons for putting fiber optics in your setup. The right reason is that there are unique physical measurements you can do with the fiber, serious safety concerns, or compelling cost advantages, and eventually lots of similar sensors are going to be made, so that it's worth putting in the engineering time to get it really right. Borderline reasons are that you need a really good spatial filter, and pinholes aren't enough, or that the system needs 100 kV of voltage isolation. The wrong reason is that your bulk optical system is inconvenient to align, and it would be nice to use fibers to steer the light around to where it's needed. If you're building sensors for others to use, you'll quite likely spend a lot of late nights in the lab, wishing you hadn't done that. Where conventional sensors (optical or not) can do the job, you should do it that way almost always.

Fiber sensors can do some really great and unique things, but they're just one of many tools for the instrument designer. If you really like fibers, and want to use them everywhere, you face the same type of problem as the man in the proverb who only has a hammer—everything looks like a nail.

8.9.1 Sensitivity

The basic trade-off in fiber sensors is sensitivity versus convenience and low cost. It's often very fruitful to go the low cost route, since the enormous dynamic range of optical measurements is not always needed, and the market potential may be enormous. Do realize, though, that you are making that trade-off, so that if it isn't low cost, there's often no point.

8.9.2 Stabilization Strategy

If we are going to make a fancy fiber sensor, we have to be able to predict its performance. Elsewhere, we've been asking "What's the SNR going to be?" and calculating it from first principles and measured device parameters. With fiber sensors, we're not that fortunate, because there are two prior questions. The first is: "How is it going to be stabilized?" Every sensitive fiber sensor is seriously unstable unless you figure out a way to stabilize it, and this instability is the primary bogey in fiber sensor work. Scale factors drift, operating points go all over the place, and all on a time scale of seconds or minutes—unless you figure out how to avoid it. Some sort of common-mode rejection scheme is usually needed, e.g. a fiber strain gauge with an unstrained reference fiber in the same package, two wavelengths, or two modes. Some of these schemes work a lot better than others.

8.9.3 Handling Excess Noise

The next question is: "How do we get rid of the excess noise?" After stability, the worst problem in laser-based fiber sensors is their huge technical noise, caused by demodulation of source phase noise by etalon fringes and scattered light. Only after dealing with these two issues do we get to the point of worrying about our SNR in the way we do elsewhere. Fiber measurements are intrinsically squishier than bulk optical ones (see Section 11.4.2).

8.9.4 Source Drift

LEDs drift to longer wavelengths with increasing T , by as much as several parts in $10^4/^\circ\text{C}$. Diode lasers drift less (at least between mode jumps) because they are constrained by their resonators. Any fiber sensor that is sensitive to changes in the source spectrum will have to neutralize these effects somehow. The easiest way is to temperature-control the source, but temperature compensation or normalization by an independent source spectrum measurement will often work too.

8.10 INTENSITY SENSORS

Most mainstream fiber sensors at present are based on changes in received intensity, at or near DC. There are many kinds, but only a few physical principles: microbending loss from pressing against corrugated plates, evanescent coupling between nearby fiber cores, and misalignment.

Intensity sensors more or less ignore phase and polarization, which makes them fairly insensitive but pretty trouble-free (for fiber devices). The limitations on their sensitivity




Figure 8.9. Microbending sensors.

come from the smallness of the signal and their vulnerability to interference from other losses, e.g. microbending outside the sensor and source drift and noise.

Intensity measurements can help us normalize the output of fiber interferometers. Sampling the beams before recombining them gives us the intensity information separately, which can be subtracted or divided out.

8.10.1 Microbend Sensors

Microbending increases loss, with the attenuation being very roughly linear in the curvature. Microbending sensors exploit this effect in a variety of ways (see Figure 8.9).

They aren't very sensitive. Most cause only a few percent change in transmission, full scale, and they tend to be fairly badly nonlinear; oddly shaped deviations of tens of percent from the best straight line are common, and that isn't too confidence-inspiring since we often don't know its origin. Microbend vibration sensors are a good technology—the AC measurement avoids most of these troubles if you're careful.

8.10.2 Fiber Pyrometers

Black body emission follows the Stefan–Boltzmann equation, so an IR fiber with a decent black body on one end and a thermal-IR detector makes a nice thermometer for high temperature applications (up to 800 °C).

8.10.3 Fluorescence Sensors

The upper-state lifetime of many fluorescence transitions goes down as T goes up. Because the decay is exponential in time, measuring the decay time normalizes out the intensity shifts. A cell full of fluorophore solution or a fluorescent coating on the fiber facet are the common ways of doing it.

An unselective detector followed by a log amp and a bandpass differentiator is all you need; this is the same basic idea as cavity ring-down spectroscopy. The most popular fiber for this is Nd:glass, whose lifetime varies from 210 μs at 0 °C to 170 μs at 250 °C, which is quite a convenient range. The repeatability of an individual, undisturbed sensor is about 0.5°, but the exact dependence of the lifetime on temperature seems to vary a lot from sensor to sensor (as much as $\pm 3^\circ$ nonlinear shift from sensor to sensor) so that for high accuracy, individual calibration over the full range appears necessary.

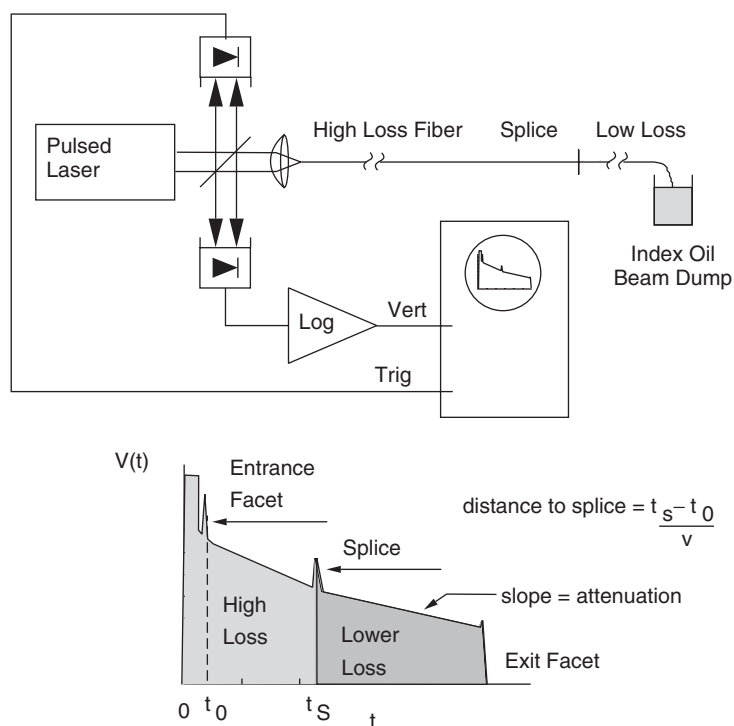


Figure 8.10. Optical time-domain reflectometry.

Fluorescence sensors based on fluorophores immobilized in fiber coatings are also used for oxygen detection; O_2 diffusing into the fiber coating quenches the fluorescence. (Of course, normal fluorescence measurements are often done using fibers as light pipes, but that isn't a fiber sensor in the sense used here.)

8.10.4 Optical Time-Domain Reflectometry

Optical time-domain reflectometry (OTDR) works a lot like radar: you send a tall, narrow pulse down the fiber, and look at what comes back, as a function of time delay. The result is a plot of $\log(r)$ versus t , that is, distance. As shown in Figure 8.10, we get a constant Rayleigh scatter signal, which decays at twice the fiber loss rate, and reflection spikes from splices and other discontinuities. A steeply dropping slope means excess loss. Chemical sensors, microbend sensors (e.g., fiber with a helically wrapped member inside the armor), and even porous cladding sensors for water can all be interrogated by OTDR. Polarization and phase are more or less ignored.

8.11 SPECTRALLY ENCODED SENSORS

8.11.1 Fiber Bragg Grating Sensors

Bragg grating sensors work by illuminating the fiber with a broadband source such as an LED, and using a spectrometer to watch the absorption peak move when the fiber

is stretched[†] or heated, in accordance with (8.24) and (8.27). Stretching the fiber moves the absorption peak by a combination of the change in length, the piezo-optic effect, and the change in the waveguide diameter through Poisson's ratio. The resolution of the measurement depends on the quality of the spectrometer and the length of the grating. This *spectral encoding* makes for a nice stable sensor, at least once the temperature dependence is removed. Fiber Bragg gratings escape nearly all the problems of other fiber sensors; they are insensitive to polarization, phase instability, excess loss, and etalon fringes (if you use fiber-coupled LEDs, you even escape alignment). They do require an external spectrometer, and the higher its resolution, the higher the sensitivity of the sensor. Because of the low étendue, you don't get that much light, but on the other hand you don't need much, because these sensors are typically measuring mechanical and thermal things, which are slow.

By putting several gratings (as many as 10 or 20 of different periods) into the same fiber at different positions, a single fiber can do simultaneous strain measurements at many positions. This idea is used especially in smart structures, which take advantage of Bragg grating sensors' ability to make measurements in otherwise totally inaccessible locations, such as the interior of reinforced concrete bridge abutments.

The strain sensitivity of this technique isn't bad; from Section 8.5.11, it's about

$$\frac{1}{\lambda_B} \frac{\partial \lambda_B}{\partial \epsilon} = 0.78, \quad (8.30)$$

so that a grating stretched by 0.1% of its length (1000 μ strain) changes its wavelength by +780 ppm. The total allowed strain isn't very big, 1000–2000 μ strain, so a sharply peaked wavelength response is obviously important. People have tried using interferometers to read out Bragg gratings, but that throws away the main advantage of the technique: its insensitivity to everything but strain and temperature. One interesting approach is to use a Mach–Zehnder delay discriminator to measure the average wavelength of the light reflected from the grating; although this is obviously highly sensitive to the exact line-shape of the reflection, it does get around the resolution limit imposed by spectrometer pixels. Fitting a grating instrument with a split detector or a lateral effect cell would be another way to do this.

The accuracy of fiber strain gauges isn't bad, a percent or so. A more serious problem is temperature drift; the tempo of λ_B is

$$\frac{1}{\lambda_B} \frac{\partial \lambda_B}{\partial T} = +6.7 \text{ ppm}/^\circ\text{C}, \quad (8.31)$$

so that a 1 $^\circ\text{C}$ change gives the same signal as a 9 μ strain stretch—which is on the order of 1% of full scale. To leading order, the two effects are both linear in λ , so separating them is difficult; the best method seems to be the one used in resistive strain gauges, that is, comparison with an unstrained grating at the same temperature. Various two-wavelength schemes have been tried, which use the fiber dispersion to make the strain and temperature coefficients different. These all fail in practice, either because the wavelengths are so far apart that the fiber is no longer single mode, or because the coefficients don't shift enough to give a well-conditioned measurement.

[†]Of course, you can do it with a tunable laser and a single detector, but then all the usual fiber problems reappear, and sufficiently widely tunable lasers don't grow on trees either.

Current laboratory accuracies are about $1\ \mu\text{strain}$ or $0.1\ ^\circ\text{C}$, but field accuracies are less because it's hard to make a good, cheap readout system. The main challenge in low resolution field systems is in maintaining a good enough bond between the strained element and the fiber that strain is transferred accurately. Quartz is very rigid stuff; glue isn't, especially not glue that's gotten hot or wet.

8.11.2 Extrinsic Fabry–Perot Sensors

Spectral encoding is also the idea behind many types of extrinsic Fabry–Perot sensors, for example, diaphragm-type pressure sensors, which look at the interference of light reflected from the fiber end and from a nearby diaphragm. Their low finesse and narrow spacing make their fringes broad enough to interrogate with white light or LEDs.

Other examples are thermometers based on refractive index drift in silicon.

8.11.3 Other Strain Sensors

Extrinsic piezo-optic strain sensors are sometimes used with fiber coupling, but this technology seems not to be becoming mainstream. Other strain-based sensors include magnetostriction-type magnetic field sensors, in which a strain sensor is bonded to a magnetostrictive element. It's pretty easy to put an unstrained companion on one of these, e.g. by bonding one with indium solder or brazing, and the other with elastomer.

8.11.4 Fiber Bundle Spectrometers

Grating spectrometers and FTIRs are big clunky things, so it's natural to want to couple them with fibers into inaccessible places. It's especially nice to use linear bundles, where the fibers are arranged in a row to fill the entrance slit. This can work OK if you pay extreme attention to the pupil function you're imposing on the measurement. An integrating sphere and a multimode bundle are stable enough for good spectra, but if you simply bounce light from one fiber off a surface and into another fiber, you're just asking for ugly measurement artifacts. You also give up a lot of light unless you use a lot of fibers; a grating spectrometer's étendue is nothing to write home about, but a fiber's is worse. It also tends to be unstable.

Two of the author's colleagues abandoned fiber-bundle spectroscopy after giving it a good run in an *in situ* application, in favor of mounting miniature monochromators right on the instrument head—they couldn't get enough photons otherwise.

One application where fibers actually *increase* the amount of light you get is in astronomical fiber plate spectroscopy, where a computer-controlled milling machine drills holes in a plate just where the stars will be imaged. Sticking multimode fibers into the plate brings all those widely separated sources into one spectrometer—a nice example of putting your étendue where you need it.

8.11.5 Raman Thermometers

The Raman effect is a bilinear mixing between vibrational levels in the fiber material and the optical field, and is similar in character to the acousto-optic interaction we talked about in Section 7.10.4. A photon can emit a phonon, and so be downshifted (Stokes

shift), or it can absorb a phonon and be upshifted (anti-Stokes). The Stokes/anti-Stokes intensity ratio depends on the occupation number of the vibrational states, and hence on the temperature. The ratio changes by a few tenths of percent/°C at a few hundred wave numbers offset, but the signal is weak; an OTDR setup with a good filter and a decent PMT or APD gives a few meters' resolution at 1 °C; by using photon counting and very narrow pulses, you can get 5 °C accuracy with 10 cm spatial resolution.

8.11.6 Band Edge Shift

Absorption and fluorescence spectra are often highly temperature dependent. The band edge of a semiconductor moves toward long wavelength with temperature, while remaining sharp, so it can be used as an extrinsic fiber thermometer. These don't get better than 1 °C, and they're almost as nonlinear as thermistors, but they can work from room temperature to 800 °C.

8.11.7 Colorimetric Sensors

Fiber chemical sensors can be made by cladding a fiber in a microporous sol-gel glass with a chromophore (color-changing chemical) mixed in. The chromophore is immobilized in tiny pores in the glass. The target chemical species diffuses in and combines with the chromophore, changing its absorption spectrum. This trick is used in commercial pH sensors, for instance.

8.12 POLARIMETRIC SENSORS

8.12.1 Faraday Effect Ammeters

Currents produce magnetic fields, which in turn produce Faraday effect polarization shifts in fibers. This underpins all-fiber Faraday rotation ammeters, in which the small Verdet constant of silica ($V = 4.6 \times 10^{-6}$ rad/A) is overcome by using many turns; with N turns wrapped around the conductor, the Faraday shift is

$$\theta = V \cdot N \cdot I. \quad (8.32)$$

This of course works only if the two circular polarizations remain distinct. If they are strongly coupled, as by bend birefringence, the Faraday effect is *quenched*; light in the fast and slow circular polarizations change places several times inside the coil, so that the phase difference tends to cancel out (a bit like twisted pair wiring). Annealing the coiled fiber can help a lot with this, or you can quench the quenching by twisting the fiber as it is coiled. The resulting circular birefringence prevents the linear birefringence from building up enough to couple the fast and slow circular polarizations much; the circular birefringence just gives rise to a fixed phase shift between polarizations, which isn't much of a problem. The fibers are relatively short, so double Rayleigh scatter isn't usually serious. The path difference between the two polarizations is small, so source PM doesn't turn into polarization jitter, and as long as the etalon fringes are well controlled, the FM-AM conversion is thus minor; the predominant noise is source intensity noise, $1/f$ noise in the front end amplifier, and the usual huge amount of low frequency junk due to temperature and bending.

In the scheme of things, nobody's about to use a fiber sensor when a better and cheaper technique such as a current transformer or a series resistor and isolation amp is applicable; thus these sensors are used only inside high voltage transmission systems. This is helpful from the noise point of view, since the current is AC, and therefore so is the measurement.

8.12.2 Birefringent Fiber

The change in birefringence with strain, pressure, and temperature is roughly proportional to the static birefringence, so people have tried using polarimetry with high birefringence fiber to sense pressure and vibration. This works at some level, but the temperature dependence is so strong, and the operating point so unstable, that such techniques appear to have little to recommend them.

8.12.3 Photonic Crystal Fiber

Recently, fibers have been developed whose cross sections are not solid but include arrays of air holes, making them 2D photonic crystals. Their structure gives these holey fibers unusual properties. Some are highly nonlinear, for example, those used in femtosecond supercontinuum comb generation; others have very wide mode fields that are very constant with wavelength, or actually guide the light inside one of the holes, rather than in the glass. They haven't been used much in fiber sensors yet, but that's likely to change soon.

8.13 FIBER INTERFEROMETERS

Fiber interferometers produce an order-unity change in output for a 1 radian phase shift, just as bulk ones do. Most of the usual kinds can be constructed in fiber: Michelson, Mach–Zehnder, Fabry–Perot, and so on.

The big challenge is that fiber interferometers are sensitive to *everything*, especially temperature. This lack of selectivity makes it hard to have much confidence in the data, unless great efforts are made. Temperature gradients are especially serious.

At the root of the problem is that the phase shift due to interfering effects tends to go as the length of the fiber; since it goes into the exponent rather than the scale factor, at some point the signal-to-noise ratio cannot be improved any more by using more fiber.

Fiber interferometers place immense demands on the stability and coherence of the laser. Ordinary diode lasers are not even in the running for this use, because their linewidths are in the tens of megahertz, and their drift, mode hopping, and mode partition noise make them very difficult to use barefoot. External-cavity stabilized diode lasers are a good choice, as are solid state devices such as diode-pumped YAGs.

The instability restricts the dynamic range of fiber interferometers. You can use them over a wide range by fringe counting, or over a narrow one by AC modulation, but you can't easily splice them together into a single wide-range measurement as you can with bulk. Even though the path length may be a factor of 1000 larger in the fiber case, fringe-counting resolution of 10^{-6} m/1 km is far poorer than the 1 Hz shot noise limit of 10^{-14} m/1 m (5 mW HeNe).

8.13.1 Single Mode

If single-mode fiber were really single mode, instead of degenerate two-mode, and that mode were really lossless, building single-mode fiber interferometers would be a piece of cake. Although neither is true, still single-mode fiber is so much superior to multimode for interferometry that it's the only game in town. We therefore make the best of it.[†]

8.13.2 Two Mode

The availability of communications fiber that supports two modes at some readily accessible wavelengths has led people to try using the mode coupling or differential delay or polarization shift between modes as a sensor. This is an admirable example of actually trying to realize the hitherto mythical cost savings due to economies of scale, instead of just talking about them. This works OK if you babysit it, but getting exactly two modes forces us to work in a regime where the second mode is somewhat leaky, and hence sensitive to manufacturing variations in the fiber, as well as to handling, bending, and so on. Higher order modes are not well preserved in splices, and especially not in connectors. It's a tough business. At present, avoid two-mode and multimode interferometric sensors if you want your gizmo to be usable by anyone but you, but stay tuned—if someone figures out how to do it well, this trick could really take off.

8.14 TWO-BEAM FIBER INTERFEROMETERS

By using fiber couplers instead of beamsplitters, it is quite straightforward to make Michelson, Sagnac, and Mach–Zehnder fiber interferometers, as shown in Figure 8.11. There are two basic approaches to using them; directly as sensors, or as delay line discriminators for optical frequency measurements. A lot of blood has been spilled making intrinsic interferometric fiber sensors, and it's just a very hard business to be in. Fiber gyros have been in development for many years and have only recently overcome their many, many second-order problems to become a mainstream technology. That's good if you're cranking out PhD theses, but not so good if you're building sensors for a living. If you're in the latter class, beware.

8.14.1 Mach–Zehnder

Mach–Zehnders are the workhorses, but they're horribly unstable if made of fiber, for all the usual reasons. Integrated optics ones are much better. All-fiber Mach–Zehnders are commonly used as delay line discriminators in high velocity Doppler shift measurements, but require a lot of babysitting due to polarization funnies.

8.14.2 Michelson

Fiber Michelsons are attractive because you can use FRMs, which eliminates the polarization drift and allows passive interrogation—the fringe visibility in a Michelson with decent FRMs normally stays above 95%, which is pretty good. The main residual problems are the few degrees of rapid polarization and phase wobble with (average)

[†]Every mode is its own interferometer, and you can't keep them all in phase.

temperature, and temperature gradients, which cause large amounts of phase drift in the relative phases of the two arms. A fiber Michelson built with FRMs and interrogated with the modulation-generated carrier approach works adequately for AC measurements, where the phase drift and fringe visibility can be taken out in software without screwing up the operating point.

8.14.3 Sagnac

Sagnac interferometers use waves propagating in the forward and reverse directions, more or less fixing the phase instability problem.

The static phase delays in the two directions are the same except for the Faraday effect, Berry's phase, and the *Sagnac effect*, which causes a phase shift in a rotating reference frame. An N -turn coil of cross-sectional area A , turning at angular velocity Ω , produces a phase shift between clockwise and counterclockwise propagating light of

$$\Delta\phi \approx \frac{8N\pi\Omega A}{\lambda c}. \quad (8.33)$$

The instability of fiber is so very well suppressed in fact that you don't get any output below a certain rotation speed—the backscattered light dominates until the Sagnac phase shift gets sufficiently large. Transient effects of course break this symmetry because they arrive at different times unless they occur exactly midway on the path.

Sagnac interferometers based on big coils of fiber have a lot of bend birefringence, just like Faraday effect ammeters of Section 8.12.1.

8.15 MULTIPLE-BEAM FIBER INTERFEROMETERS

8.15.1 Fabry–Perot

There are two kinds of fiber Fabry–Perots: *intrinsic* ones, where the fiber is entirely within the cavity, and *extrinsic* ones, where at least part of the cavity is outside the fiber (Figure 8.11).

An intrinsic F-P is a stretch of fiber with a Bragg grating or a mirrored facet at each end. It is sensitive to all the usual F-P things, minus misalignment, plus the usual fiber things, which is a really bad trade. Polarization instability prevents the use of intrinsic F-Ps in unattended situations and severely limits the attainable finesse; if you don't mind babysitting, however, their potentially very long path length makes them extremely sensitive detectors of tuning, strain, and temperature, *provided* that your laser linewidth is narrow enough. Because the cavity is normally long, the free spectral range is short; a 10 m piece of fiber has an FSR of 10 MHz, and if its finesse is 50, the FWHM of the passband is 200 kHz. Diode lasers typically have linewidths of 10–100 MHz, so the fringes show up in the spectrum rather than in the intensity, making all-fiber F-Ps hard to use except with very fancy lasers.

Low finesse intrinsic F-Ps can be made by cleaving the desired length, depositing a TiO_2 film on each end, and fusion-splicing the result back into the main fiber. Reflectivities are 1–10%. Because of their low finesse, they are optically much more like unbalanced Michelsons or Mach–Zehnders, but need no couplers. Another trick is the fiber loop mirror of Figure 8.11c, in which a cross-connected fiber coupler sends part

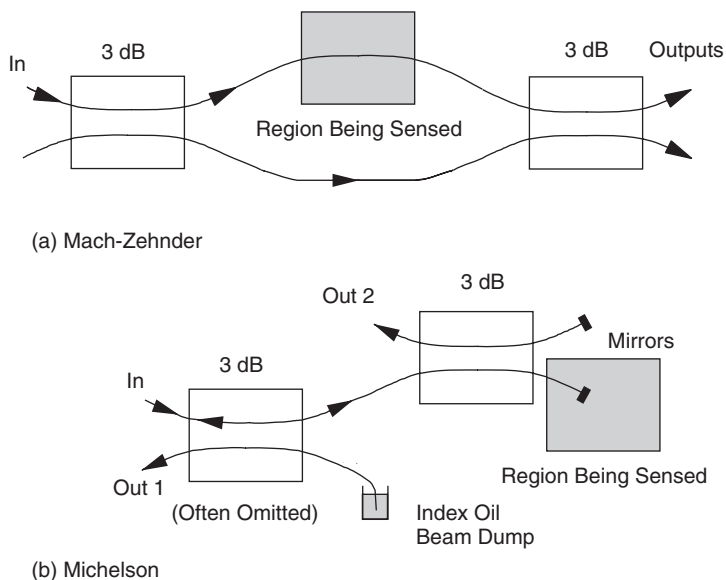


Figure 8.11. Two-beam fiber interferometers.

of the light back on itself. If the power coupling ratio is $x : (1 - x)$ and we neglect depolarization, the reflectance is

$$R = x(1 - x) \sin^2(\phi). \quad (8.34)$$

This simple fiber device can be used as a two-beam Sagnac interferometer itself, or as a mirror in a larger system such as a fiber laser. Since the reflectance is a very strong function of wavelength, temperature, and polarization, some sort of tuning will probably be required.

Extrinsic F-Ps come in two flavors: fiber delay lines with one or more bulk optic mirrors, and fiber coupled F-Ps, where the fiber just brings light to the real Fabry–Perot. An extrinsic F-P with one Faraday rotator mirror corrects the polarization problems of the original, and in the process makes the effective length of the cavity polarization independent, thereby removing one source of ambiguity. (See Figure 8.12.)

8.15.2 Ring Resonator

A ring resonator is very similar to an F-P, except that its transmission looks like the reflection of an F-P. It has all the same problems, too. If you launch light going both ways, you can use a ring as a resonant fiber-optic gyro. A good technique to eliminate polarization problems in a ring resonator is to use PM fiber, but splice it rotated 90° , so that the two polarizations exchange identities each turn. This actually makes the periodicity twice as long, but produces two bumps per period; adjacent passbands are slightly different.




Figure 8.12. Fiber Fabry-Perot sensors: (a) extrinsic, with external diaphragm; (b) intrinsic, with multilayer coating spliced into fiber; and (c) fiber loop (Sagnac interferometer used as a mirror).

8.15.3 Motion Sensors

Doppler-type motion sensors are a good match for fiber interferometry; you put a collimator on the output of one fiber of a Michelson interferometer and point it at the thing you want to measure. Absolute phase stability is not usually needed, and providing the polarization instability is controlled, all the low frequency headaches can be filtered out, and fringe-counting accuracy is usually good enough. Near-field applications such as fiber catheter flow cytometry don't need the collimator, but far-field ones do, and a bit of retroreflecting tape is usually a big help in getting light back into the single-mode fiber.

If the radial velocity is too high to count fringes ($\mathbf{v} \cdot \mathbf{k}/2\pi > 100$ MHz or so), you can use the fiber interferometer as a delay discriminator on the returned light instead.

8.15.4 Coherence-Domain Techniques

There is a family of techniques for measuring displacement and distributed parameters such as temperature without needing a fast photon-limited OTDR setup (Figure 8.13). They rely on the idea of white-light fringes. The measured output of an interferometer is a short-time-averaged autocorrelation between the signal and LO, which means that for a very wideband source fringes appear only when the path lengths of the two arms are very close. The key feature of coherence-domain measurements is that, apart from polarization funnies, the autocorrelation of the signal is preserved pretty well through the fiber. Some dispersion exists, of course, but that is compensated by the equal-length reference arm. Wideband sources such as LEDs have only a single autocorrelation peak,

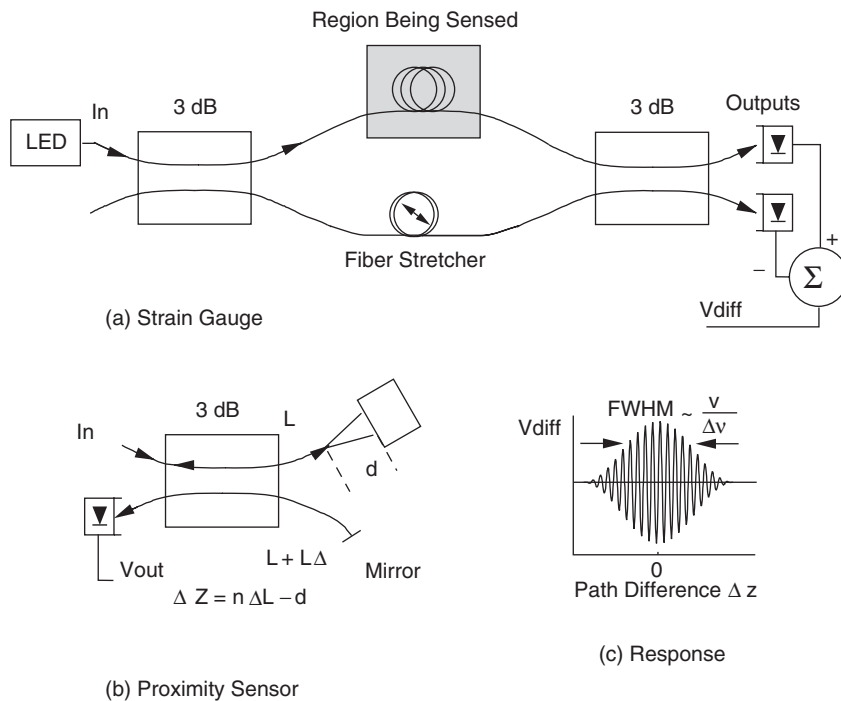


Figure 8.13. Coherence-domain techniques: (a) strain gauge, (b) OCDR proximity sensor, and (c) typical response curve.

unlike Fabry–Perots and highly coherent interferometers. Thus by stretching the reference arm until the white-light fringes appear, we can measure the distance to a scatterer unambiguously.

The accuracy is correspondingly less, because we have to locate the center of the autocorrelation peak, which of course has a fast carrier frequency imposed on it; thus unless we can use an optical I/Q mixer to get the true envelope, there is always a one-fringe error band.

Because the autocorrelations of multimode lasers tend to be about 10^3 fringes wide, the SNR needed to get one-fringe accuracy is far higher than with a fringe-counting interferometer: 50–60 dB instead of 15 dB, which is a serious limitation. An LED with a 5% bandwidth is rather better, but on the other hand, you don't get as much light. One way to get around it is to synthesize an autocorrelation function that is much less smooth, and *in a known way*. For example, we can use a slightly unbalanced Mach–Zehnder to put fringes on the autocorrelation, as shown. This can also be done by using two wavelengths, or a filter with two passbands on a single wideband source. It is analogous to a dual-wavelength interferometer, where the fast fringes give high resolution and the long beat length disambiguates the 2π phase wraps. This is an example of a multiple-scale measurement, which we discuss in Section 10.4.2. Coherence-domain techniques are good for low resolution length measurements, such as noncontact surface sensing, for interrogation of features in fiber that are far too close for OTDR (e.g., 1 mm apart), and for disambiguation of interferometric data. They're also good for measurements in turbid or multiple-scattering media such as tissue.

8.16 PHASE AND POLARIZATION STABILIZATION

Getting a stable operating point in a fiber sensor is a bit of a challenge, as we've seen. There are lots of methods in use, of which the following are the best.

8.16.1 Passive Interrogation

If the paths are sufficiently unequal (>1 cm), the phase and amplitude can be interrogated separately using a tunable laser. To avoid the linewidth problems we had with the Fabry–Perot, the laser should be tunable over many times its linewidth, and the nominal path difference adjusted accordingly—for example, $\Delta L = 1.5 - 2.5$ cm for a CD-type diode laser whose tuning range is 1 cm^{-1} . This allows stabilizing the operating point electronically, by tuning the laser.

One good way to do it is with the modulation-generated carrier approach, as in Section 13.9.6. This technique requires modulating the phase of the light by ± 2.6 radians or so, where the J_1 and J_2 terms in the Bessel expansion of the phase modulation are equal. Because you can combine in-phase and quadrature (I and Q) signals so as to ignore the slow phase drift completely (perhaps with a feedback loop to get the balance exactly right), this is a pretty effective solution to the phase problem. You do have to keep the modulation stable at that value, and the SNR suffers surprisingly little, since 85% of the interference term energy is in the first and second harmonics.

8.16.2 Frequency Modulation

Sufficiently rapid FM modulation can help a bit in reducing the coherence fluctuation noise, by transforming the widely different time delays of the twice-scattered components into large frequency shifts so that most of that contribution lies outside the measurement bandwidth. The modulation frequency needs to be several times the laser linewidth, and the modulation index should be large. Still, as we saw in Section 4.7.4, this technique isn't as good as we'd like. Very rapid chirps turn distance into frequency shifts, which is often better than simple FM.

8.16.3 Fringe Surfing

You can use a fiber stretcher or other phase modulator to surf on a fringe, essentially phase-locking to the fringe drift. This scheme, discussed in Section 10.7.7, is the first method everybody thinks of. Unfortunately, the phase will in general walk a long way over long times, far further than phase modulators can accommodate; for instance, a YAG laser resonator that changes by 0.001% in length will change the delay of a 50 m piece of fiber by 0.001% of 70 million wavelengths, or 1400π radians, and a 1.5°C temperature change in the fiber will do the same. Thus fringe surfing works only in the lab unless your measurement can be interrupted for a second or two at irregular intervals to reacquire lock. Because of the limited bandwidth of most such setups (fiber stretchers resonate in the hundreds of hertz to low kilohertz), they are seriously microphonic, too—bump the table and the loop will lose lock, to come back to rest somewhere else tens of milliseconds later.

As with a PLL phase demodulator (Section 13.9.5), you can take as the system output either the detected output from the optical phase detector, before the loop integrator, or the control voltage to the fiber stretcher, after the integrator. Unfortunately, the control

voltage here is useless for measurements. There are multiple operating points on the fiber stretcher, so every time the loop resets, there will be a huge voltage step on that output; furthermore, the actuator is nonlinear, so there is no easy way of reconnecting a run of data that has been interrupted by a reacquisition.

8.16.4 Broadband Light

Using broadband light and spectrally encoded data makes a lot of things easier, provided the source drift is taken care of. The main drawback is the pixel resolution of the spectrometers used to decode the spectrum.

8.16.5 Ratiometric Operation

One way to get rid of the scale factor drift is by comparing the signal to a known one; for example, a Faraday effect gaussmeter could use an electromagnet to excite the coil stably at a known frequency; the ratio of the two signals would correspond to the ratio of the applied field to the standard one. This requires only high linearity, and not high stability. Avoid trying to use two wavelengths, because etalon fringes among other things are highly tuning sensitive.

8.16.6 Polarization-Insensitive Sensors

Polarization drift can be greatly reduced by using Faraday rotator mirrors, as we've seen, and although this isn't magic, it keeps the polarization drift well enough controlled that the operating point doesn't move around much. The residual drift will still need correcting for in many cases, but that can be done in postprocessing.

One source of residual polarization drift is polarization-dependent losses, for example, Fresnel reflection at angled fiber cleaves. These sorts of effects produce errors by coupling the supposedly orthogonal modes together, thus ruining the orthogonality on which the FRM trick depends. Another is light that doesn't get transformed in the FRM, for example, facet reflections and front surface reflections, and multiple reflections in the Faraday crystal. A third is errors of the FRM itself, for example, misadjustment, dispersion, drift, and nonuniformity.

8.16.7 Polarization Diversity

The modulation-generated carrier approach is an example of phase diversity; it is also possible to use polarization diversity, although this is much more complicated. One method is to combine the ideas of Sections 6.10.10 and 15.5.4; after recollimating the output light, use bulk optics to split it into two or three copies. Use polarizing prisms rotated through different angles to select different projections of the two fields, and pick whichever one has the best fringe visibility. This takes care of the problem of misaligned major axes, leaving only the usual relative phase problems.

8.16.8 Temperature Compensation

Separating temperature from other effects usually involves measuring temperature independently and using a calibration curve. Another property of the fiber (e.g., Raman

backscatter) can be used, or an IC temperature sensor. The particular problems with this are generating the calibration curve, and figuring out what temperature to use, in the face of significant gradients across the dimension of the fiber system.

8.16.9 Annealing

You can anneal fiber that has been wound round a mandrel by heating it up to 800 °C for a while and cooling it slowly, being sure not to stretch it by letting the mandrel expand too much. That gets rid of the bend birefringence, but it tends to be hard on the jacket, which makes it unattractive unless you really, really need to, e.g. in Faraday sensors for power transmission ammeters. It is wise to leave the fiber undisturbed on the mandrel afterwards.

8.17 MULTIPLEXING AND SMART STRUCTURES

Some types of intrinsic fiber sensor lend themselves to multiplexing, which permits the use of many sensors with only one readout system, and perhaps with only one interrogating fiber. Multiplexing can be done in all the standard electronic ways; time-division (TDM), frequency-division (FDM), and code-division (CDM), plus a uniquely optical one, *coherence multiplexing*. The details of this are really beyond our scope, being network and interconnection issues, but they are discussed in Udd and elsewhere.

8.18 FIBER SENSOR HYPE

As the old saw runs, “It’s not what you don’t know that hurts you, it’s what you do know that ain’t so.” Fiber sensors are currently still fashionable, though not as modish as they were. In part, this is because of the advantages already enumerated, although as we’ve already seen, there’s less than meets the eye about some of them. The author is somewhat sceptical of fiber sensor claims in general, because the sheer volume of hype obscures any clear-eyed assessment of their strengths and weaknesses, and because the discussion in fiber sensor papers—even review papers and book chapters—seldom seems to include comparisons to bulk optic alternatives. We’re now in a position to enumerate and critique a few.

1. *Fiber Sensors Offer Unique Capability.* Some types certainly do. Fiber gyroscopes, Bragg grating sensors, distributed sensors like Raman thermometers, and arguably Faraday effect current sensors all have capabilities that are difficult or impossible to replicate with bulk optics. Fiber-optic smart structures, oil-well thermometers, and fiber catheter devices all have an important place.

2. *Fiber Sensors Are Cheap.* They could be, in principle, but sensitive ones almost never are—people who are used to the prices of electronic sensors are liable to keel over from pure sticker shock. Some intensity-based systems, such as fluorescent fiber thermometers, really are intrinsically cheap. Those tend to be the ones where the fiber is just used to pipe the light around, which is fair enough. Distributed systems, such as fiber Bragg grating strain gauge systems, where the cost of the interrogation system is

amortized across hundreds of sensors, might one day be cheap on a per-sensor basis, which is why they are interesting for fiber-optic smart structures.

Cheapness is often said to be a consequence of economies of scale in the telecommunications business, but it ain't so. Telecommunications fibers are not single mode at the wavelengths of cheap sources and detectors, and the cost of fiber transmission systems is not driven by the cost of lasers and detectors, as fiber sensor costs are. Besides, the *real* economies of scale are in diode laser-based consumer products—which are 100% bulk optics. (Compare the cost of a 1.5 μm diode laser to a 670 nm one of the same power level, if you doubt this.)

Simplicity and low cost can rapidly go away when we start dealing with phase and polarization instability, etalon fringes, FM–AM conversion, and extremely low étendue; we often find ourselves using \$10,000 single-frequency diode-pumped YAG lasers to get decent performance from \$50 worth of fiber.

It's also common to find apples-to-oranges comparisons; if a sensor has to work in situ, compare a fiber system to a miniaturized and cost-reduced bulk optic system, not to some behemoth with an argon laser and a table full of Newport mounts. Think in terms of DVD drives and digital cameras.

3. *Fibre Sensors Are Highly Sensitive.* This is sometimes true compared with nonoptical sensors, but is more or less entirely false in comparison to bulk optics. Leaving aside the expensive kinds, like gyros and Faraday rotator mirror interferometers, the really unique fiber sensors are all relatively low sensitivity devices. What's more, the performance of even fancy fiber sensors falls far short of the fundamental physical limits.[†] Fibers are highly sensitive to temperature, bending, vibration, you name it, and that leads to enormous amounts of low frequency drift and noise. High frequency performance is limited by the multiplicative effects of the low frequency stuff, including drift of the scale factors and operating points—things we'd never tolerate in a DVM—and by demodulation of source phase noise by etalon fringes and double Rayleigh scattering.

There are an awful lot of interferometric fiber sensor papers where somebody is bragging about his 120 dB dynamic range in 1 Hz. That isn't a bad number, but in reading these, remember that a bulk optic interferometer with a \$5 diode laser and a \$10 laser noise canceler just sits there at the shot noise level forever, unattended, with a dynamic range of 160 dB in 1 Hz; and that performance isn't hype, it's available in commercial hardware.[‡] Don't forget to ask about their relative stability and $1/f$ noise, either.

[†]This isn't entirely a fiber problem, of course.

[‡]To be fair, getting the noise canceler down to \$10 requires building your own—the commercial ones are over \$1000. The good part of this is that it's easy to do.