

Interplay Between Physics and Statistics for Modeling Optical Fiber Bandwidth

Steven Golowich
Mathematics Research
Bell Laboratories
Murray Hill, NJ 07974
(golowich@bell-labs.com)

James Landwehr
Data Analysis Research
Avaya Labs
Basking Ridge, NJ 07920
(jml@avaya.com)

Scott Vander Wiel
Statistics Research
Bell Laboratories
Murray Hill, NJ 07974
(scottyv@bell-labs.com)

Abstract

In multi-mode optical fiber the refractive index of the glass is varied radially in such a way that a light pulse propagates through multiple modes, or paths, as it travels. We investigate predicting fiber transmission capacity, specifically bandwidth, by using refractive index profiles of glass preform rods measured in a manufacturing environment before the rods are drawn into optical fiber. By closely linking empirical and theoretical approaches to modeling, we demonstrate the feasibility of predicting bandwidth in spite of the finding that profile measurements are grossly inaccurate. Empirical components of the modeling involve extensive pre-processing of raw measurements that approximate continuous functions, followed by spline fitting and specialized cross-validation for model assessment. Modeling from physical theory is based on solutions of Maxwell's equations. This paper shows how a blending of the physics with the statistics provides conclusions that could not be obtained separately from either approach—namely, good predictions are possible but profiling equipment needs to be upgraded to remove substantial measurement biases.

Keywords: Data Analysis, Empirical Model, Physical Model

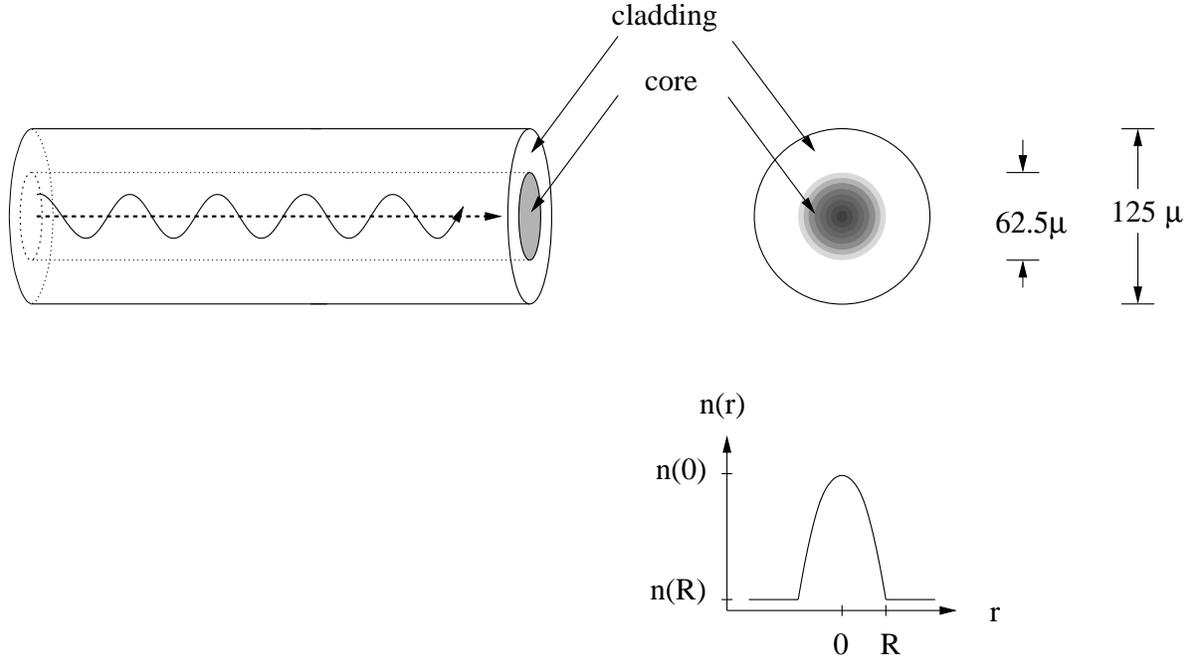


Figure 1: Schematic of rays (dash: straight; solid: oscillating) traveling through a graded index fiber. The refractive index of the core is graded to equalize travel times.

1 Introduction

The need for higher bit rates in local area networks has increased the demand for a certain type of optical fiber known as multi-mode fiber. A simplified view of light propagation through multi-mode fiber is shown in Figure 1. The fiber consists of a 62.5μ (micron) glass core surrounded by a glass cladding, and the light travels down the core. The *refractive index* of glass is inversely proportional to the speed with which light travels through the glass. Refractive index gradients built into optical fiber determine, via Snell's Law, how light bends as it travels through the fiber core. In multi-mode fiber the refractive index is graded with the radius in such a way that multiple modes—or, roughly speaking, multiple paths—of light propagation are supported. We refer to the function $n(r)$ that describes the variation of index with radius as the *profile* of the fiber. Figure 1 depicts two rays traveling through a graded index fiber and the general shape of the refractive index profile, $n(r)$.

We would like to predict fiber bandwidths from measured refractive index profiles within the context of a manufacturing environment. A wealth of well-known physical theory specifies how bandwidth relates to refractive index in the ideal case. This theory, which is based

on Maxwell's equations, is summarized briefly in Section 2 though many details of the physics are omitted. Standard refractive index profile designs are based on the theory, but they also involve assumptions and simplifications to the real world such as axial symmetry and longitudinal homogeneity of the refractive index and specific knowledge of material properties of the glass. In practice the assumptions hold well enough that theoretically desirable refractive index profiles, when manufactured properly, do provide adequate bandwidth.

We investigate the extent to which measured *departures* from the ideal refractive index profile can be related through physical theory and empirical models to measured *differences* in bandwidth among fibers. While theoretically such a relationship must exist, practically speaking it is not obvious. Are measurements and simplifying assumptions accurate enough so that typical manufacturing departures from the ideal fiber correspond in a reproducible, understandable, and measurable way to differences in bandwidths across fibers?

Relating refractive index profiles to bandwidths is important for two main reasons. If the refractive index can be measured accurately enough early in the production process, manufacturing costs could be reduced: bad product could be discarded early on, and the production process could be adjusted to increase bandwidth. Second, if the adequacy of certain theoretical approximations can be justified empirically, then it would be possible to use the theory with confidence for more complex system design problems and for modeling light propagation in fiber used under more stringent operating conditions.

This paper illustrates the process of modeling in a manufacturing setting that leads to two results: (1) demonstration that propagation properties in multi-mode fiber can be predicted from refractive index profiles as measured in either preform or fiber cross-sections; and (2) indication that the profilers used to measure preforms have bad accuracy and need upgrading in order to be used for bandwidth prediction. These results are achieved through a close linkage between empirical and theoretical approaches to modeling. Section 3 describes flexible empirical modeling used to establish that preform profiles do have predictive information about bandwidth. Popular opinion had been that preforms are too difficult to measure with enough accuracy for bandwidth prediction. While the empirical model demonstrates predictive ability its cross-validated multiple correlation of 0.63 is ultimately not good enough to reliably sort out bad preforms or tune the production process.

The empirical work is driven by good physical understanding but it does not bring in

theory in a fundamental way. Section 4 describes a modeling approach that is primarily grounded in the theory but requires empirical adjustments to be successful. This model predicts propagation delays of the individual transmission modes that carry light down the fiber. The mode delay pattern theoretically determines bandwidth, and is fundamentally important to understanding fiber performance. The predicted delay pattern based on preform profiles is much too widely dispersed to be consistent with measured bandwidths, which suggests that the preform profiling equipment needs to be upgraded. Two different types of *fiber* profile measurements produce a narrower, more realistic delay pattern, even though fiber profiles are intrinsically more difficult to measure than preforms due to the smaller physical scale of fibers. Furthermore, after some empirical adjustment, delay variation predicts bandwidth to about the same accuracy as obtained with the empirical model. These results are reasonable in terms of the underlying physics and convinced manufacturing engineers that improvements to the preform profiling equipment could produce a big payoff: namely, measurements that are sufficiently accurate to use for tuning the production process and for removing bad preforms early in the product stream.

In addition to developing and assessing the two models, a major purpose of this article is to highlight the qualitative differences and advantages of the two approaches to model building in the context of a difficult engineering and manufacturing problem. Such issues are addressed through the paper and revisited in Section 5 where we show that improvements to the quality of preform profile measurements do indeed allow accurate and straightforward prediction of light propagation through fiber drawn from the preforms. We emphasize the importance of bringing empirical thinking to bear on physical model building, and, conversely, of blending physical insight into empirical modeling.

2 Background

2.1 Refractive Index Profiles and Mode Propagation

Glass fiber transmits digital data from a light source, such as a laser or an LED, that is switched on and off. The source launches light into one end of the fiber so that a series of pulses is sent down the fiber. If successive pulses stay well separated in time and do not attenuate too much, the signal is transmitted. A *mode* of light propagation is defined more

precisely below, but we may think of a mode as a specific propagation path inside the fiber: examples are the path straight along the axis, an oscillatory path crossing the axis (as in Figure 1) and a path winding helically around the axis. When a pulse of light enters the fiber, it propagates as a superposition of many pulses, each traveling in a different mode and carrying a fraction of the total power. If different modes propagate at different speeds, the “pulse” spreads out as it travels. If successive pulses are sent too close together in time they will blur into a single wide pulse as they travel. This blurring effect determines the bandwidth that a fiber can support.

The refractive index profile $n(r)$ specifies the value of the refractive index as a function of the radial distance r from the center of the core and determines which modes will propagate and how fast each travels. Fiber with a refractive index profile that varies from a high index at the center of the core to a lower index at the core-cladding boundary, as shown in Figure 1, is called graded-index fiber. The refractive index grading is designed so that all propagating modes travel at approximately the same speed, thus preventing excessive spreading of a composite light pulse as it travels down the fiber with power distributed across many modes. Below is a heuristic explanation of why this works.

Consider simultaneously launching two very small spots of light into the end of a fiber and parallel to the axis but at different cross-sectional positions. One spot is launched into the center of the core and, according to geometric optics, travels undeflected straight down the center of the fiber (e.g., dash path in Figure 1). The other spot is launched at some radial distance r from the center of the core and it will oscillate back and forth between $\pm r$, passing repeatedly through the center as it travels (e.g., solid path in Figure 1). The oscillating path is obviously longer than the straight path. The transit time of the straight path, however, may not be shorter, because that pulse spends all of its time in the “slow” glass with higher refractive index. The other pulse travels through glass that has a lower refractive index on average, so the speed is higher and compensates for the longer path. If the refractive index profile is chosen correctly, the two pulses will exit the far end of the fiber simultaneously.

One might worry that there are other paths, or modes, with different transit times even if we equalize the above pair. It can be shown (Snyder and Love, 1983) that all modes have approximately equal transit times if the index profile within the core is chosen from the

so-called α -family with properly chosen α . For this family the refractive index profile $n(r)$ at radial distance r from the center of the core is given by

$$n^2(r) = n_0^2 \left(1 - 2\Delta \left(\frac{r}{R} \right)^\alpha \right), \quad 0 \leq r \leq R \quad (1)$$

where n_0 is the index at the center of the core; R is the core radius; and 2Δ , which measures the fractional difference between the squared-index of the center of the core and the cladding, is typically around 0.02. In practice the optimal profile within this family depends on the material properties of the doped silica used to manufacture the fiber. Optimal values of α are typically very close to 2, meaning that a parabolic refractive index profile is nearly ideal.

We typically measure refractive index as a percentage above the refractive index of the pure silica found in the cladding region outside the fiber core. Thus, we define

$$\begin{aligned} \text{Refractive Index}(r) &\equiv \frac{n(r) - n(R)}{n(R)} \\ &= \left(\frac{n^2(r) - n^2(R)}{2n_0^2} \right) \left(\frac{n_0^2}{n^2(R)} \right) \left(\frac{2n(R)}{n(r) + n(R)} \right) \end{aligned}$$

where the second equality comes from algebraic manipulation. As $n^2(r)$ varies by less than 2% over r , the final two factors are nearly unity. Using this approximation and then substituting Equation (1) into the first factor, we see that for the α -profile family

$$\text{Refractive Index}(r) \approx \Delta \left(1 - \left(\frac{r}{R} \right)^\alpha \right), \quad 0 \leq r \leq R. \quad (2)$$

Thus, Δ is approximately the fractional index difference between the center of the core and the cladding.

The preceding description has identified modes with paths that obey the equations of geometric optics and gives a useful mental picture. A more accurate description is that light is a wave phenomenon governed by Maxwell's equations. For our purposes we do not need to discuss the details of Maxwell's equations and their solutions for this problem, but we briefly describe relevant characteristics of the solution. A *mode* represents a solution to an eigenvalue problem from a wave equation. The distinguishing feature of a propagating mode in an ideal waveguide is that, when excited, it will propagate indefinitely without losing energy to radiation or other modes. If one follows the wavefronts of a given mode, they follow paths similar to those described in the geometric optics picture of propagation given above. However, only a finite set of allowed modes solve the equation, not the continuum

of possible paths that was suggested from the geometric viewpoint. A standard multi-mode waveguide supports several hundred modes. They can be divided into *principal mode groups* (PMGs), of which there are typically 20 to 30. In real fibers the individual modes within a PMG tend to share power freely among themselves because of micro-scale variations that are always present due to manufacturing processes. Therefore, we may think of a PMG as one large, composite mode.

Summarizing, we can think of optical fiber as having a refractive index profile such that light is propagated among 20 to 30 modes. A light pulse has power distributed across these modes. Dispersion of the pulse results because modes may travel at different speeds along the fiber.

A real-world optical fiber offers additional complicating factors, of course. A traveling light pulse loses power from absorption, from light scattering outside the fiber due to material impurities, and from conversion to radiation due to bends in the fiber. The amount of loss varies from mode to mode. While the ideal theory assumes perfect axial symmetry in the refractive index profile, in practice non-negligible asymmetries may exist and can affect mode propagation. Material properties of the glass influence mode structure and propagation speeds. Although the relevant properties can be measured, the degree of uncertainty is not negligible. Finally, the preceding discussion is all conditional on the wavelength of the light source. Everything changes as a function of wavelength, and standard multi-mode fiber must have good properties for two wavelengths, 850 nm and 1300 nm.

Manufacturing optical fiber primarily involves two distinct processes. First, a *preform* rod is built. The preform is a solid glass cylinder, with the inner part corresponding to the core of the fiber and the exterior annulus corresponding to the cladding. Typical preform sizes are diameters in the 10's of millimeters, and typical lengths are in the meters. Each preform is built over many hours on a lathe where the delivery of several chemicals is carefully controlled so that the preform has the desired purity and refractive index properties. After some characterization and post-processing, the finished preform is taken to a draw tower, which can be several stories high. There the tip of the preform is heated and the glass is drawn continuously and pulled into a thin strand of optical fiber 125μ in diameter. The fiber is immediately coated and wound on a spool. A single preform can produce tens of kilometers of fiber.

In principle, the draw process should not substantially alter the refractive index profile. That is, aside from the diameter change in going from the preform to the fiber, the refractive index profiles should be otherwise practically unchanged.

3 Empirical Modeling

3.1 Preform and Fiber Measurements

Refractive index profiles can be measured on both preforms and fibers. For *preforms*, a laser light source is moved through a cross-section of the preform and deflection angles of the light are recorded as the light exits the cross-section; deflection angles are determined by how much of the core and cladding the laser beam passes through and from the refractive index of this glass. The light passes sideways through a cross-section of the preform, not along the preform axis. Deflection angles are input into a proprietary mathematical inversion routine within a commercial measurement tool in order to produce the refractive index profile for one cross-section of a preform.

Routine refractive index profiles from *fiber* are obtained by a method based on the same underlying physics, but with a different measurement process since the fiber dimensions are too small to move a “point” light source through a fiber cross-section. Without going into details, these are called transverse interferometric measurements of refractive near-field index profiles from fiber.

Figure 2 shows an example of refractive index profile measurements from a preform. The refractive index values are available at approximately 1400 points equally spaced across a diameter of the preform core, with the center labeled as 0.0 and the left and right radii of the core labeled $-R_l$ and R_r respectively. We refer to each half of the profile as a *trace*; from 0 to $-R_l$, or from 0 to R_r . We often work with each trace separately. Several features are apparent in this typical profile measurement. There is an overall approximately quadratic shape with refractive index near 2.0% at the center and declining to near 0 at the edges of the core. This corresponds to the theoretically desirable shape of Equation (2) with $\alpha = 2$ and $\Delta = 2\%$. Small cycles, or ripples, throughout the trace vary in amplitude and radial width. The fact that such ripples exist is known to result from the manufacturing process, and to some extent the ripples are unavoidable. A deep dip appears near 0, and a near-linear

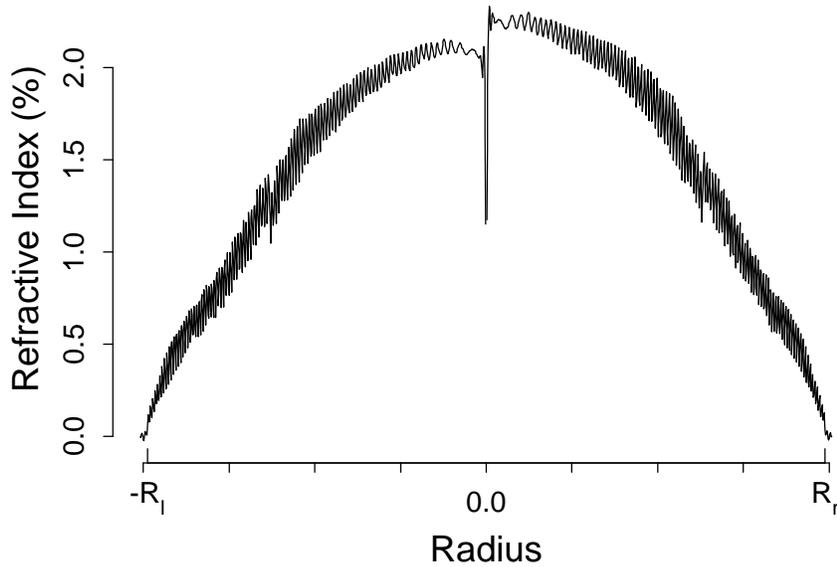


Figure 2: The core section of a typical refractive index profile as measured through a cross-section of a glass preform.

section appears in the middle of each trace where a true quadratic would remain curved. The left and right traces are similar but not perfectly symmetric. The two tick marks above the horizontal scale indicate the core radii for the left and right traces. Although nearly equal in this measurement, the left and right radii are not generally equal. They typically differ by less than a few percent.

This section describes a modeling approach emphasizing an empirical, statistical perspective using refractive index measurements on preforms such as shown in Figure 2 to predict transmission capacity as measured by fiber *bandwidth*. Bandwidth is commonly given as the frequency at which the Fourier transform of the impulse response first drops to one-half its peak value. Alternatively, bandwidth can be expressed through directly measuring the root-mean-square width of the impulse response and applying the formula

$$\text{BW (MHz km)} = \frac{170}{\text{RMS pulse width (ns/km)}} \quad (3)$$

which has been shown empirically to be a very good approximation (Buckler, 1982). Ideally bandwidth is inversely proportional to the length of fiber. Thus, bandwidth is usually multiplied by length and is reported in units of MHzkm. It is not possible to measure

bandwidth directly on a preform because the measurement requires sending a signal down a fiber.

While a substantial body of physical theory exists concerning relationships between refractive index profiles and optical fiber transmission, this theory had not been “operationalized” for use with realistic features of typical measurements so as to provide a theoretically-based modeling approach that could be used for bandwidth prediction in a manufacturing environment. A successful empirical approach, though perhaps not ideal from the point of view of physical understanding and extrapolation, would provide valuable steps towards resolving important engineering problems.

Empirical bandwidth modeling proceeds as follows: refractive index measurements are pre-processed and B-spline approximations are used to represent the resulting explanatory functions. The B-spline degrees of freedom are set by cross-validation to choose an appropriate level of complexity and avoid over-fitting. Alternatives to B-splines for flexibly modeling a response from functional predictors are discussed by Kiiveri (1992) and Ramsay and Silverman (1997).

Our pre-processing procedure removes types of profile variation that are unimportant according to fiber optics theory. This is analogous to curve registration common in functional data analysis (Ramsay and Silverman, 1997). Reducing such measurement artifacts can be crucial for successful empirical modeling especially with the large space of relevant models for making predictions from curves. Reasons for ignoring various artifacts are discussed to highlight the interplay between physics and statistics. A nice side-effect of the pre-processing is that the extensive raw profile data as shown in Figure 2 are reduced many fold to a more manageable size and structure for further analysis. The reduced profiles are more readily visualized alongside of bandwidth measurements (see Figure 5) which is extremely helpful in guiding the empirical modeling.

3.2 Theoretically Informed Pre-processing

Available data consist of refractive index profiles measured at two longitudinal positions on each of 184 preforms. The preforms were produced on a set of 13 lathes. Each preform was drawn into several spools of fiber and bandwidth was measured on each spool. At the time of this study it was not possible to know the correspondence between a particular spool of

fiber and the longitudinal segment of the preform from which it was drawn. Even though the cross-sectional profile and fiber bandwidth may change longitudinally this effect cannot be modeled with available data; for this reason we characterize a preform by one “overall” profile and a corresponding “overall” bandwidth. Further studies will attempt to incorporate the effects of longitudinal changes on bandwidth.

Figures 3 and 4 illustrate the pre-processing of data for one preform. Refractive index profiles measured at two positions along the length of the preform are shown vertically offset in the top panel. Ideally the preform is axially symmetric and longitudinally homogeneous so that the four traces (half profiles) at any two positions would be equal. These traces, while visually similar, are clearly not identical. The differences are likely due both to inhomogeneity in the actual preform and to measurement problems, but the amount of each is unknown.

The first pre-processing step is to fit the radius R for each trace by finding the location of a small jump in the refractive index where the core and cladding meet. The details of how we do this are not important but tick marks above the horizontal axis indicate the 4 radii. Next we choose to eliminate, for both theoretical and practical reasons, two portions of each trace: the central portion with $r/R < 0.08$, and the cladding portion with $r/R > 1$. We omit the cladding portion because it is essentially constant and therefore uninteresting. Near 0 the manufacturing process induces certain glass characteristics that can affect light transmission in ways that affect refractive index measurements but not fiber bandwidths; so the profile measurements near 0 do not relate to the part of the problem we are attacking. Furthermore, from observing many profiles it appears that measurements are substantially more variable near 0 than elsewhere in the trace. One reason is that the commercial algorithm that converts laser deflection angles into profiles may sometimes have trouble determining the exact center of the preform core. Mis-locating the core badly biases the profile near 0. Thus, for several reasons, the central 8% of each trace is removed in preprocessing. Typical variations in core radius are known to have little effect on bandwidth (Olshansky, 1979). For this reason profile measurements are simplified by rescaling to a unit radius—that is, the radial values on the horizontal axis are divided by the measured core radius R so the rescaled values run from zero to one.

The solid curves drawn on each trace in the top panel of Figure 3 are baseline profiles

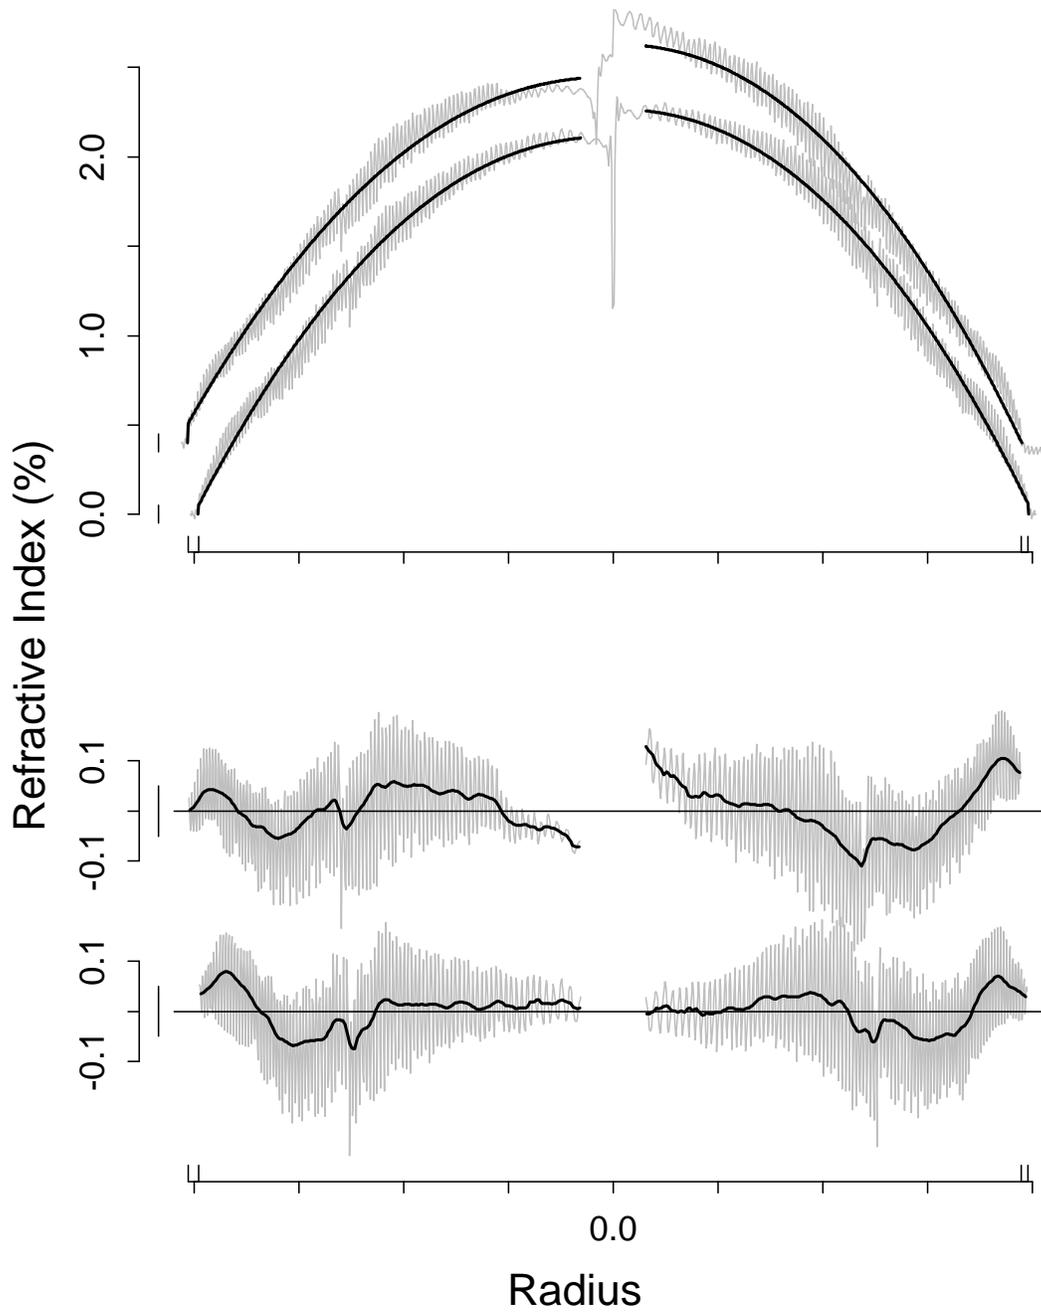


Figure 3: Fitting and smoothing of the profiles from a single preform. The top panel shows original measurements from two cross-sections with fitted ideal profiles. The bottom panel shows residuals from the fitted ideals and smooths of those residuals.

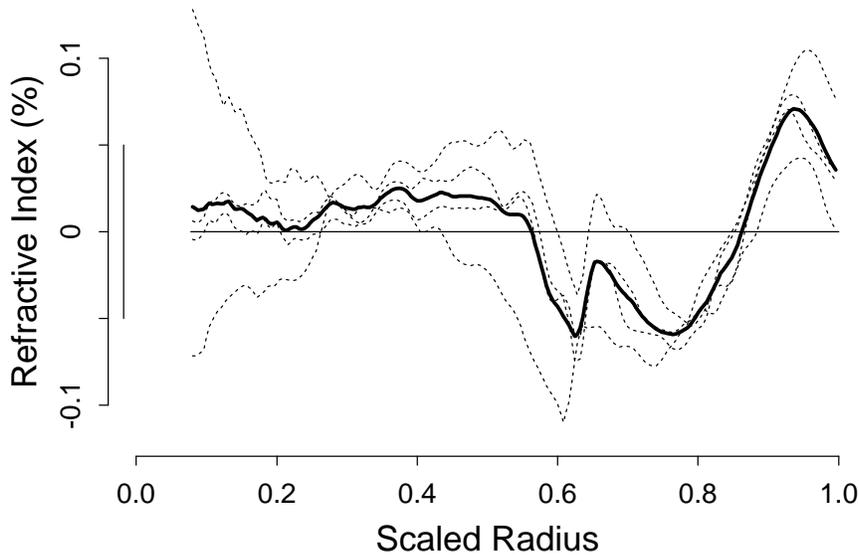


Figure 4: Four smoothed profiles (dotted) and their average (solid). The average curve contains the profile information used for modeling bandwidth.

obtained as least squares fits of Δ_0 and Δ to

$$\text{Refractive Index}(r) \sim \Delta_0 + \Delta \left(1 - \left(\frac{r}{R}\right)^2\right), \quad 0.08 \leq r/R \leq 1. \quad (4)$$

This is Equation (2) with an added intercept parameter, Δ_0 , and $\alpha = 2$. Recall that $\alpha = 2$ is the theoretically desired value for maximizing bandwidth. We call Δ_0 the *jump* because it measures the small refractive index jump at the core to clad interface.

The lower panel of Figure 3 shows the result of subtracting the theoretical baseline curve from each of the four traces and omitting the data near zero and beyond the measured core radius. Removing the baseline structure allows finer visual comparisons among the traces. Vertical line segments of length 0.1% are drawn to the left of each profile in Figures 3 and 4; from there, one can appreciate how much variation is removed at different stages of the pre-processing.

A smooth local regression curve is drawn through each trace in Figure 4. A near discontinuity at approximately half the core radius is now easily visible in the processed measurements, and, as later determined, stems from an experimental step in the preform processing in use during this phase of product development. Discontinuities appear as kinks in the smooth curves. The span for the smoother was chosen wide enough so that the ripples

are removed yet the trends remain. Removing the ripples obviously brings out the larger-scale structure and available physical theory suggests that the small-scale ripple structure should have little effect on bandwidth. Also, in manufacturing it is easier to control the smooth structure than the ripples. We did, however, attempt to use some measures of ripple amplitude, in addition to the smooths, when building a bandwidth model.

Δ_0 and Δ type differences have also been removed in the pre-processing and these, too, are expected to have minimal impact on fiber bandwidth. This can be justified both theoretically and from engineering considerations. We do, however, record the fitted values of Δ_0 and Δ and consider these as potential predictor variables for modeling bandwidth.

In Figure 4 all four smooths are plotted with standardized unit radii. The left-side traces from Figure 3 have been reflected about zero. The final step of the preprocessing is to interpolate the 4 smooths onto a common scaled radius grid (231 values from 0.08 to 1.00) and calculate the average smooth, which is shown as the dark solid line in Figure 4. There are two justifications for reducing the four traces to a single average trace. First, averaging greatly reduces the kinds of measurement effects that arise when the equipment has trouble determining the exact center of the core; see the two traces in Figure 4 that diverge near radius 0 but have little effect on the average. A second justification for averaging comes from optical waveguide theory by which a longitudinal and rotational average profile is the most important determinant of fiber bandwidth.

The preprocessing described above was applied to refractive index profiles from each of 184 preforms. Table 1 indicates the remaining variation after each pre-processing step as well as the reduction in variability attributed to each step. “Profile Mean Square” is the average sum of squares along all profiles at the given stage of pre-processing. In constructing the table, the initial pre-processing step of fitting and removing separate baselines from each individual trace is split into two steps. First a common baseline (equal to the best fit of Equation (4) to all traces) is removed, and then separate fits to each trace are further removed. Thus the profile variation from Δ_0 and Δ -type *differences* among the profiles can be quantified.

The second column in Table 1 lists the differences in mean squares between pre-processing steps. Thus, subtracting multiple baselines reduces profile variations by 5.31, which is more than 13 times the amount of variation remaining in the final pre-processed profiles (0.40).

<i>Pre-processing Step</i>	<i>Profile Mean Square</i> ($\times 1000$)	<i>Reduction in</i> <i>Mean Square</i>	<i>Percent</i> <i>Reduction</i>
Original measurements	2141.70		
Common baseline removed	10.88	2130.82	99.5%
Multiple baselines removed	5.57	5.31	48.8
Smoothed	0.68	4.89	87.8
Averages of 4 traces	0.40	0.28	41.2

Table 1: Profile variation remaining after each step of pre-processing and amounts removed by each step.

Smoothing accounts for almost as much reduction. The final pre-processed profiles that we use for bandwidth prediction represent less than 4% ($= 0.40/10.88$) of the original profile variation about a common baseline. Furthermore, as shown in the third column, each step removes a large fraction of the variation from the previous step. If, as we expect, the pre-processing retains the profile information relevant to predicting bandwidth, then letting the optical fiber theory guide the pre-processing has enabled us to greatly increase the signal-to-noise ratio for the next step of building an empirical model.

3.3 Bandwidth Predictions from B-Spline Models

Figure 5 shows a measure of bandwidth (left panel) and corresponding pre-processed profile measurements (right panel) for all 184 preforms produced on 13 lathes over a 3 week time period. We want to predict bandwidths from the profiles. Kinks from the experimental processing step mentioned previously are visible in the profiles from each lathe.

For the response variable, each preform is drawn into several spools of fiber and bandwidths are measured. We use the median log bandwidth as the response to be predicted, because other parts of the manufacturing process can occasionally have problems that result in abnormally low bandwidth fiber for reasons separate from those being investigated here and the median over spools is robust to such effects. Further, note that within lathe bandwidth variation is not strongly related to the mean in the left panel of Figure 5, where the bandwidths are plotted on a log scale.

Manufacturing engineers know that different lathes can produce different quality fiber for reasons that are not entirely understood, and that the performance of lathes can improve or degrade over time. A premise of our modeling is that differences in fiber quality may result from lathe to lathe differences in typical profiles. Hence, one use of a bandwidth model would be to support tuning lathes to to produce preforms with better profiles. For this reason and because theory says that transmission properties should be determined by the profile shape, we do not use the lathe identification as a bandwidth predictor for purposes of model development. We do, however, incorporate lathe information into model assessment.

Our modeling approach is to capture important variations among the pre-processed refractive index profiles to predict bandwidth. Because the space of all profiles is large and we have bandwidths from only 184 preforms, we need to impose smoothness in a prediction model that uses the profiles. We use an approach that Ramsay and Silverman (1997) describe as *regularization using basis functions*. Let \mathbf{X} be a matrix with 184 rows containing the pre-processed profiles:

$$\mathbf{X}^T = (\mathbf{x}_1, \dots, \mathbf{x}_{184})$$

where \mathbf{x}_i is a vector containing the profile for preform i evaluated at 231 scaled radial positions from 0.08 to 1.00. Then we fit a model

$$\mathbf{y} \sim \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\boldsymbol{\delta}$$

where \mathbf{y} is a vector of log-bandwidths, and the matrix \mathbf{Z} has columns corresponding to an intercept and possibly other predictor variables not included in the pre-processed profiles. In the example discussed below, \mathbf{Z} has columns containing the values of Δ_0 and Δ that were used to remove the baseline shape given by Equation (4).

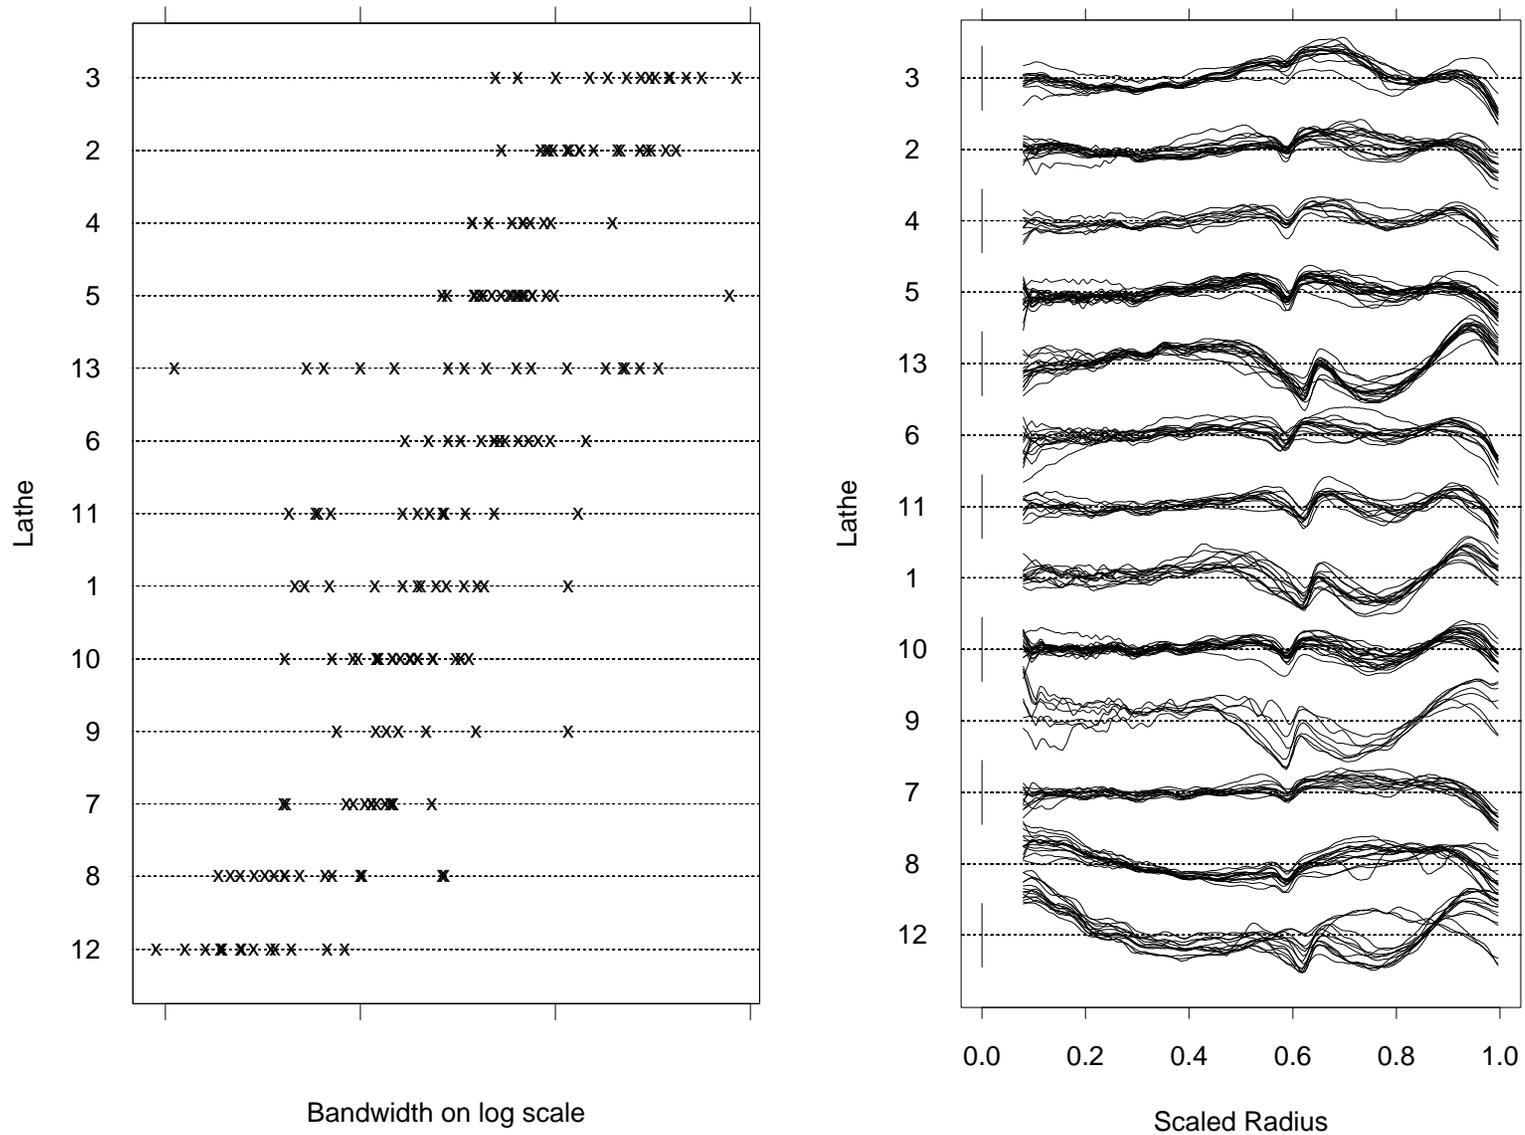


Figure 5: Bandwidth measurements (left) and processed profiles (right) for 184 preforms from 13 lathes.

Removing the baseline from each trace implies that

$$\mathbf{X}(\mathbf{1}, \mathbf{r}^2) \approx \mathbf{0}$$

where $\mathbf{1}$ is a vector of 1's and \mathbf{r}^2 contains squares of the 231 scaled radial positions where the profiles \mathbf{x}_i are evaluated. (The equation is approximate only because the subsequent smoothing step used to remove the ripples does not preserve the orthogonality exactly.) This restriction on \mathbf{X} means that we want to restrict $\boldsymbol{\beta}$ to be orthogonal to $(\mathbf{1}, \mathbf{r}^2)$ because the components of $\boldsymbol{\beta}$ in that subspace will be arbitrary and uninterpretable. Beyond this orthogonalization, it is critical to impose some smoothness on $\boldsymbol{\beta}$ to avoid over-fitting. We do this by restricting $\boldsymbol{\beta}$ to a k -dimensional subspace (with k to be determined):

$$\boldsymbol{\beta} = \mathbf{B}_k \boldsymbol{\gamma}$$

where

$$\mathbf{B}_k = (\mathbf{b}_1, \dots, \mathbf{b}_k)$$

and the basis vectors \mathbf{b}_j are residuals from regressing natural cubic spline basis functions (Schumaker, 1981) onto $(\mathbf{1}, \mathbf{r}^2)$. For k basis vectors we define the natural splines using $k + 1$ knots equally spaced from 0.08 to 1. This is the default basis given by the `ns()` function in S (Becker, Chambers and Wilks, 1988).

In summary, we predict log-bandwidths by using least squares to produce a prediction equation

$$\hat{\mathbf{y}} = \mathbf{X}\mathbf{B}_k \hat{\boldsymbol{\gamma}} + \mathbf{Z}\hat{\boldsymbol{\delta}}. \quad (5)$$

Another way to view the regularization introduced by \mathbf{B}_k is that it reduces the complexity of information supplied by the profiles in \mathbf{X} . This follows from the fact that

$$\mathbf{X}\mathbf{B}_k = \hat{\mathbf{X}}\mathbf{B}_k$$

where $\hat{\mathbf{X}}^T$ is the projection of \mathbf{X}^T onto \mathbf{B}_k . Thus we can view model (5) as using only the approximate profiles in $\hat{\mathbf{X}}$ for predicting \mathbf{y} . Figure 6 compares actual and approximate profiles for two preforms for dimensions $k = 2, 4,$ and 8 . As k increases, more of the fine-scale features of the profiles \mathbf{x}_i are captured by the approximations.

As k increases the number of basis functions used to predict bandwidth using Equation (5) increases, and the overall goodness-of-fit generally increases. (This is not strictly

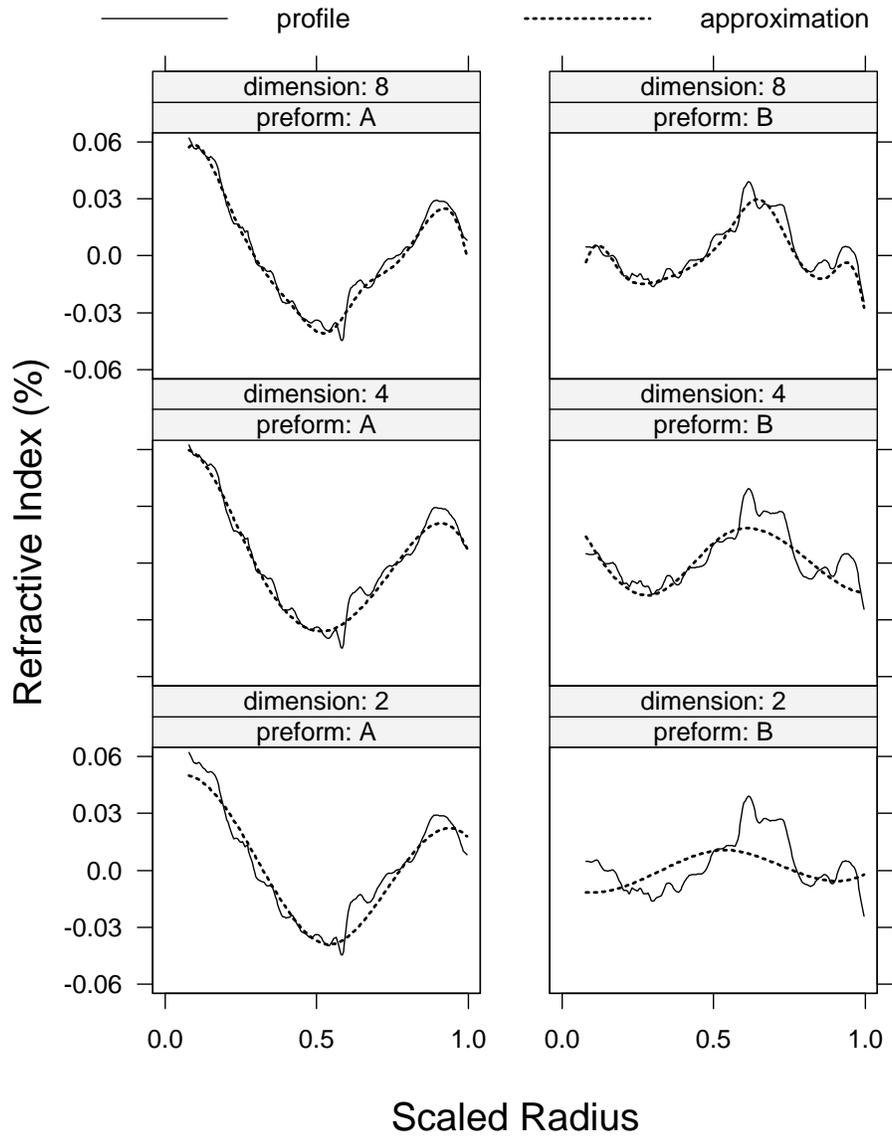


Figure 6: Successive approximations for two pre-processed profiles using spline bases. Higher dimension approximations capture more features of the profiles.

true, however, because our choice of knot selection implies that the column-space of \mathbf{B}_k does not contain that of \mathbf{B}_{k-1} .) The solid curve in Figure 7 shows the calculated correlation \hat{r} between \mathbf{y} and $\hat{\mathbf{y}}$; generally it increases with k . An important issue, however, is the choice of k . Fine-scale features of the traces should be used only if they represent meaningful signals related to bandwidth and not just noise.

We determine k using the following cross-validation procedure. The 184 preforms are divided randomly into 13 sets of about 14 preforms each. (Section 3.4 gives the reason for using 13 sets.) One set of preforms is set aside; we fit Equation (5) to the remaining preforms and use this model to predict log bandwidth for each preform in the held-out set. This process is repeated for all 13 sets giving predictions \hat{y}_i for each preform and a corresponding cross-validated multiple correlation of $\hat{r} = \text{cor}(\hat{y}_i, y_i)$. The middle line in Figure 7 shows values of \hat{r} for each k . This curve generally increases for k from 1 to 6 and declines slightly for larger k . This suggests up to approximately 6 basis functions capture structure in the preform profiles that is verifiably related to log bandwidth through this modeling process, but more than 6 basis functions does not really help prediction. These additional basis functions may be capturing mainly noise. Consequently, we tentatively adopt model (5) with $k = 6$, which gives a cross-validated \hat{r} of 0.71. While this is a long way from perfect prediction of log bandwidth, it is also substantially larger than zero and would provide useful engineering results if it performed similarly on future product.

3.4 Model Assessment

Statistical techniques such as various types of residual plots, leverage calculations, and specialized goodness of fit calculations can be useful for model assessment. Rather than explore these techniques, however, this section compares the cross-validation used for model building with another type of cross-validation that incorporates lathe information. Our goal is to assess whether the empirical model satisfactorily captures lathe effects.

Manufacturing engineers know that some lathes produce preforms whose fiber has higher bandwidth than other lathes, for reasons that are neither completely understood nor entirely under the control of the engineers. According to physical theory, the transmission properties should be almost completely determined by the refractive index profile, which is built into the preform while it is processed on a particular lathe and the profile is fixed after the

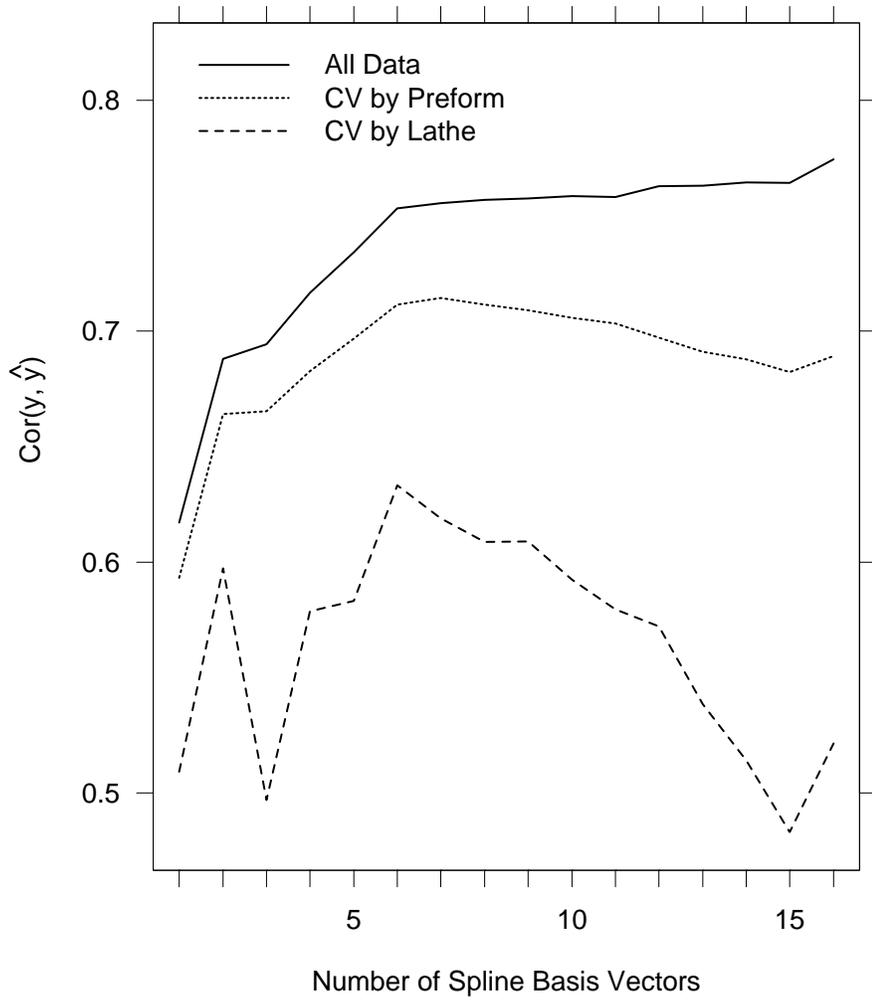


Figure 7: Multiple correlation from fits using all data and two types of cross-validation. Cross-validating over lathes performs worse than equivalent 13-fold cross-validating using random preforms.

preform leaves the lathe. Figure 5 demonstrates that lathes differ substantially in both their typical profiles and their typical bandwidth distributions. A premise of our modeling is that, although bandwidths differ from lathe to lathe, these differences stem only from profile differences among the lathes. If the premise is correct, then the model should generalize to preforms constructed on other lathes or the same lathes a year from now as long as the profiles do not fall too far outside the envelope of profiles used for developing the empirical model. In the manufacturing environment the 184 preforms come from only a subset of lathes in the factory, so a satisfactory model should also be applicable for preforms from additional lathes for which we have no comparable data.

Lathe-wise cross-validation is used to assess the premise that lathe identification provides no *additional* predictive power *beyond* our modeling of the preform profiles as follows. Leaving out all the preforms from one lathe, we fit the model (5) to all preforms from the remaining 12 lathes, use this model to predict log bandwidth for each held-out preform, and calculate residuals and the cross-validated multiple correlation. The values for each k are plotted as the lower curve in Figure 7. This curve peaks at $k = 6$ with value 0.63, again suggesting the use of 6 basis functions in the model. The correlations are also substantially larger than 0 indicating that the model has predictive power for preforms from new lathes. Nevertheless, the lathe-wise cross-validated correlations are substantially smaller than the corresponding values from the previous preform-wise cross-validations. Leaving out all preforms for a given lathe and developing the model using the other lathes gives substantially poorer predictive power compared with leaving out preforms at random. The implication is that while the model might generalize to new preforms from these 13 lathes as suggested by the middle curve in Figure 7, the model does not appear to generalize nearly as well to preforms from additional lathes. This also suggests that the model might not generalize well for these 13 lathes as their characteristics change over time. We conclude that our model using only refractive index measurements does not capture all of the differences among lathes that are relevant to bandwidth. The empirical model would likely prove inadequate for the complete engineering goals if used over time on new preforms and new lathes.

We can raise several possible reasons that lathe identification could have predictive value beyond the profile model expressed in Equation (5), such as the following: the pre-processing may have removed meaningful signal; the B-spline approach might not adequately capture

the signal; measurements of the preform profiles might be inaccurate; the profile might need to be measured on a finer grid; and information within the preforms, yet separate from the profile, may affect transmission properties and differ on average from one lathe to another. Several of these possibilities could be investigated using a more intensive empirical approach, yet along lines similar to the approach of this section. Instead, we chose a modeling strategy that attempts to incorporate the physical theory in a more fundamental way. The new model (see Section 4.2) is based on optical physics and indicates that measured preform profiles are grossly distorted relative to the accuracy needed to predict bandwidth. Refining the empirical model built in this Section could not have uncovered this fundamental problem. Further comparisons between empirical and physical approaches to modeling are discussed in Section 5.

4 Modeling Based on Physical Theory

4.1 Overview and Assumptions

A modeling approach that directly follows the underlying physical theory has two major components. First, given a refractive index profile, computer code that solves Maxwell's equations calculates the relative mode delays. The code we used, based on the finite element method, is described by Lenahan (1983). The physical characteristics of the glass medium affect these calculations through a set of constants that must be incorporated into the code. These constants represent 12 coefficients from a standard Sellmeier expansion, and have been fit to measurements by many authors, for example Fleming (1978). The result of the calculations, based on input of a specific refractive index profile and a specific set of Sellmeier coefficients, can be displayed in a differential mode delay plot as illustrated in the right panel of Figure 8. The horizontal axis indicates the principal mode group (PMG) index and the vertical axis displays the average differential delay across all modes within each PMG. An ideal situation with maximum bandwidth would have all modes with equal delay.

Proceeding from the differential mode delay structure to bandwidth (or, using (3), the RMS pulse width) involves calculations on how the power is distributed among the PMG's. That is, if some modes have delays quite different from other modes but do not carry much power, the differential delays for these modes would not have a large detrimental effect on

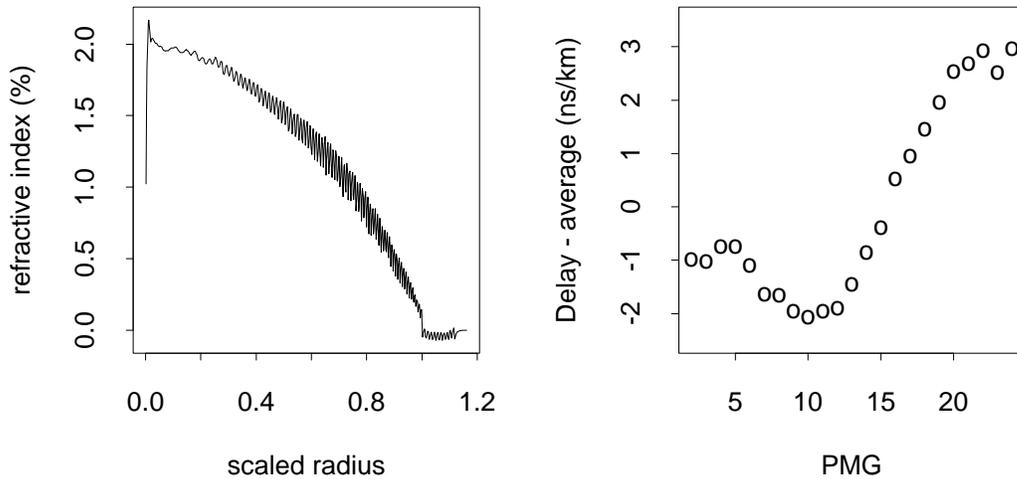


Figure 8: An example of a modal delay calculation. The left panel displays a measured profile. The right panel shows the predicted modal delays.

bandwidth. Thus, the second conceptual part of this physical modeling approach involves assumptions or calculations concerning power distribution among modes to go from the differential mode delay plot to bandwidth.

As with any theory, the above approach involves several simplifying assumptions. The accuracy and usefulness of this approach in practice will depend on the extent to which the assumptions and approximations are met, even though the underlying theory is well accepted. The following list highlights the major areas of assumptions and uncertainties, which we will address in the rest of the section.

- The code that solves Maxwell’s equations is one-dimensional, in the sense that the refractive index profile is assumed to be cylindrically symmetric. To relax this assumption would have been computationally unwieldy.
- Differential modal attenuation (DMA). If all of the high-order modes, which typically have the largest differential delays, are heavily attenuated, then in principle the bandwidth due to the remaining modes can be high.

- Mode mixing. In principle, if the PMG's shared power among themselves rapidly enough, the bandwidth of the fiber could be much larger than theoretically predicted through differential mode delay plots.
- Uncertainty in material properties. Different sets of Sellmeier coefficients have been published in the literature. Different sets could give different mode delay structures, and choosing the appropriate set is difficult and can substantially affect the results.
- Sampling and measurement error. Conceivably our measurement grid may not be fine enough to adequately specify the refractive index profile. The magnitude and structure of the measurement error, which are not known directly, could be large relative to the resolution and accuracy needed for the theoretical calculations to be useful.
- Draw effects. We are concerned with refractive index profiles from both preforms and fiber. To a first approximation these should be the same; that is, the preform profile is not modified during the draw process, only the radial scale is changed. However, in actuality the detailed conditions during the draw, such as temperature, tension, and speed, could have an impact. Obtaining and modeling data from such conditions are beyond the scope of our efforts, so such effects, if large, could affect the applicability of the theory to both preform profiles and fiber profiles.

The remainder of this section describes how we have used available data to apply a physically-based model, assess its accuracy, and determine its usefulness for our purposes.

4.2 Using the Physical Model to Evaluate Preform Measurement Accuracy

Figure 8 displays in the right panel the differential mode delay plot calculated from the preform profile shown in the left of the figure. The most salient feature is that the high-order modes have delays significantly larger than those of the low-order modes. The spread in delays is greater than 5 ns/km. Fibers drawn from this preform had measured bandwidths in excess of 1000 MHz km. According to (3), however, such a bandwidth can have an RMS pulse width of at most 0.17 ns/km. Thus, a large inconsistency, on the order of a factor of 25, exists between the modeled mode delay structure and the mode delays based on

bandwidth measurements. Similar results hold for other preforms for which we compared these calculations to measured bandwidths.

Possible explanations for this discrepancy are that the measurements, the theoretical calculations, or both, are flawed. We consider ways in which the physical model could give inaccurate results and make the following assessments.

Concerning the uncertainty in material properties, we obtained nine different sets of measured Sellmeier coefficients which might be applicable to the preforms in our study. These parameter sets induced noticeable differences in the delay structures but could not bring the model predictions in line with measurements, since the differences between the delays induced by the parameter sets were far smaller than required. If this collection of parameter sets covers a range not far from the true values for the glass in our fibers, then uncertainty in material properties could not cause a discrepancy as large as observed.

Concerning differential mode attenuation, some standard measurements known as differential mode delay (DMD) were performed on certain fibers. In these measurements only selected modes are excited and then observed at the end of the fiber. All DMD measurements indicate that many of the high-order modes do, in fact, propagate with little attenuation. Thus, differential mode attenuation cannot explain the huge discrepancy.

Similarly, mode coupling is unlikely to be sufficient to account for the discrepancy, in that any fiber with such large coupling would also have very high loss due to coupling to radiation modes. However, our fibers had very low loss. Furthermore, the result of a heavily mode-coupled fiber would be a nearly Gaussian impulse response, whereas typical fibers can have very non-Gaussian, even multimodal, impulse responses.

Another possibility is that of draw effects. Standard multimode fibers are typically drawn under a wide range of conditions with no known, reproducible effects on bandwidth. It seems exceedingly unlikely that a draw effect could change the profile as much as would be required to alter a poorly performing refractive index profile into a very good one, while being independent of such a range of draw conditions.

The modal delay calculations were run on a variety of possible profiles to examine possible effects of the pre-processing. For example, calculations were done including ripples and on corresponding profiles with ripples smoothed away, and on profiles before and after averaging longitudinally and across both traces.

Our conclusion is that none of the areas of uncertainty in the physical model could plausibly account for the size of discrepancies reported at the beginning of this section. Therefore, there are likely substantial inaccuracies in the preform refractive index profile measurements. The preceding evidence and discussion suggest that they are more likely to be biases than simply large, symmetric random error. It is impossible to uncover the existence of a large, systematic bias from an empirical analysis such as in Section 3. Moreover, a large measurement replication study would provide information only on the magnitude and structure of the random variation part of measurement error but not biases.

An obvious question is whether the bias is consistent enough in relationship to measurement variability so that we can correct empirically in some way for the bias through post-processing. We tried various ideas along this line without success and concluded that the bias is not sufficiently consistent. In the next section, we consider different types of profile measurements for which we *were* able to correct observed biases and ultimately make quantitative bandwidth predictions.

4.3 Evaluation of Fiber Profile Data by the Physical Model

This section and the next analyze profiles taken directly on *fiber*, as opposed to the *preform* profiles discussed previously. The primary motivation for examining fiber profiles is to provide a proof of the concept of predicting bandwidth from measured profiles by way of the physical model. There is no inherent reason why preform profiles should be less accurate than fiber profiles, so success in bandwidth prediction using fiber profiles provides strong motivation to upgrade the preform measurement equipment, remove measurement biases, and predict the bandwidth built into the glass before it is even drawn into fiber.

We first describe a use of the physical model that validates certain fiber profile measurements by comparing measurements of the same profile using two very different techniques. We claim that the physical model provides a useful “metric” by which to compare the different measurements, even in the face of the various model uncertainties described in the previous section. A direct comparison of two index profiles (e.g., using L_2 or L_1 norms) is difficult to interpret, because certain types of noise and biases can cause two profiles to appear very different while having almost no effect on the resulting modal delays. On the other hand, certain very small profile deviations can result in large changes in the delays. Applica-

tion of the physical model converts raw measured profiles into predicted modal delays, which are the relevant quantities for comparison.

Fiber profile measurements were available on 32 fibers drawn from seven different preforms. The measurements include the bandwidth of each fiber and two sets of index profiles taken using the Transverse Interferometric (TI) and Refractive Near Field (RNF) methods, which we call the interferometric and near field sets, respectively. These two measurement techniques are discussed by Marcuse (1981) but for present purposes it suffices to note that the measurements operate on different principles and thus can be expected to exhibit different biases.

The interferometric measurements are more routine in the factory and were obtained for all 32 spools of fiber. The fiber on the outside end of each spool was measured at both left and right radii and at both 0 and 90 degree orientations. Because little fiber is typically lost between spools drawn from the same preform, the profile from the outside end of one spool must be the same as the profile from the inside end of the next spool. Using this correspondence, we have fiber profile measurements from both ends for all but the first fiber spool drawn from each preform.

The near field measurements were taken in a laboratory setting, which rendered them both more trustworthy and much more expensive. For this reason only a subset of three fibers from the interferometric set were remeasured along four evenly spaced diameters using the near field technique.

Initial inspection of the raw interferometric profiles reveals the need for significant preprocessing to render them useful. The most serious problem is a consistent difference between the left- and right-hand sides of the profiles, which occurs to some degree in almost all of the profiles at both angles. A typical example is illustrated in Figure 9. The profiles are consistently larger on the right side than the left side, even though the fibers were measured at two uniformly spaced diameters and were loaded into the measurement device in a random orientation. This kind of difference can be attributed only to measurement bias. Identifying the cause was not possible, because the profile measurements are the result of a complicated inverse problem and the algorithm was not under our control.

A second problem with the interferometric data is that large perturbations are often present at radii $r < 7\mu$. Most fibers have some perturbations in the index profile at the

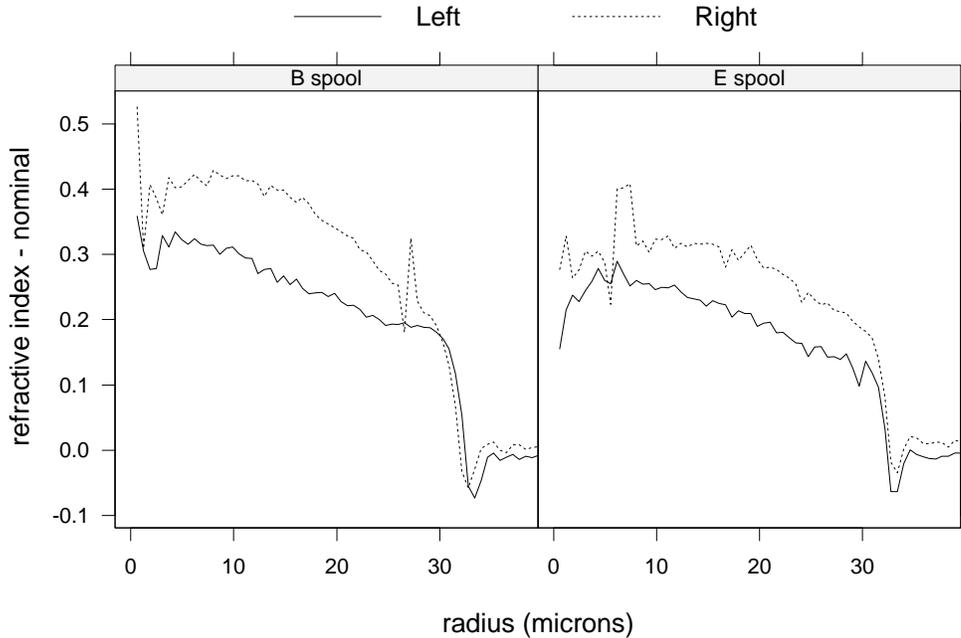


Figure 9: Two interferometric fiber profiles. The noise and biases seen in these examples are typical.

center, but in a number of examples an off-center depression is apparent that is too large to be plausible. Both the left-right bias and large perturbation problems might arise from the way the profile reconstruction algorithm chooses the center of the fiber, since the reconstruction is sensitive to the choice of center due to a singularity in the inverse problem at $r = 0$. The algorithm assumes rotational invariance, making the observed biases potentially serious. A third problem is a significant amount of noise throughout the profiles, including occasional large spikes believed to be the result of debris in the instrument.

We attempted to correct the biases through preprocessing similar to that applied to the preform profiles. This correction involved the following three steps. First, the raw profiles were smoothed to eliminate most of the short range noise and spikes. Second, averages were performed over the profiles taken along all four radii at each fiber end. While it is not clear how accurately this computed average represents the underlying true average profile over these orientations, it is plausible that the unknown sources of the biases could cancel, at least to first order. Third, the large perturbations near the center were eliminated by replacing the profile values in the central region $r < 7\mu$ by the values from the corresponding α -profile

that was fit to the smoothed, averaged measurements from outside this central region. This preprocessing represents fairly drastic measures to make this set of interferometric profiles potentially usable. As measurement processes improve in the future, we expect that such severe pre-processing would not be required.

With this preprocessing we are able to compare the interferometric profiles to those from the near field method to evaluate the extent to which the former may contain useful information. To do this, we also calculated angular averages over the near field profiles and applied the mode delay calculations to both the processed interferometric and near field profiles of the same fiber. The results are shown in Figure 10 for the three fibers. Comparing the first two rows, the two sets of delay vs. mode plots are very similar for each fiber. Moreover, the third row displays the differences and suggests that much of the differences can be explained by a single linear function of mode. While the differences in the third row do not all fall exactly on the common straight line – and the departures from linearity for the right half of the modes show similarities across the three fibers – the agreement is remarkable. Because near field and interferometric are very different measurement techniques, this correspondence leads us to conclude that the measurements in both sets do, in fact, contain useful data. While we cannot explain the differences, we suspect that the origin is related to the cause of the left-right differences in the interferometric profiles that was discussed above.

4.4 An Empirically-Adjusted, Physical Theory Model to Predict Bandwidth

This section applies adjustments to the physical model reflecting departures of real fiber from the theory to produce quantitative bandwidth predictions in three steps. The first two steps are preprocessing, by smoothing and averaging similar to Section 3.2, and computing the modal delays averaged over the PMGs. The third step involves post-processing the computed delays by modifying them with the results from an empirical model. These modified delays are then used to predict the RMS width of the impulse response, and hence the fiber bandwidth.

We discussed the areas of uncertainty for real fiber in comparison to the physical theory in Section 4.1. With our sparse and limited (yet expensive) data on 32 fibers, we must limit

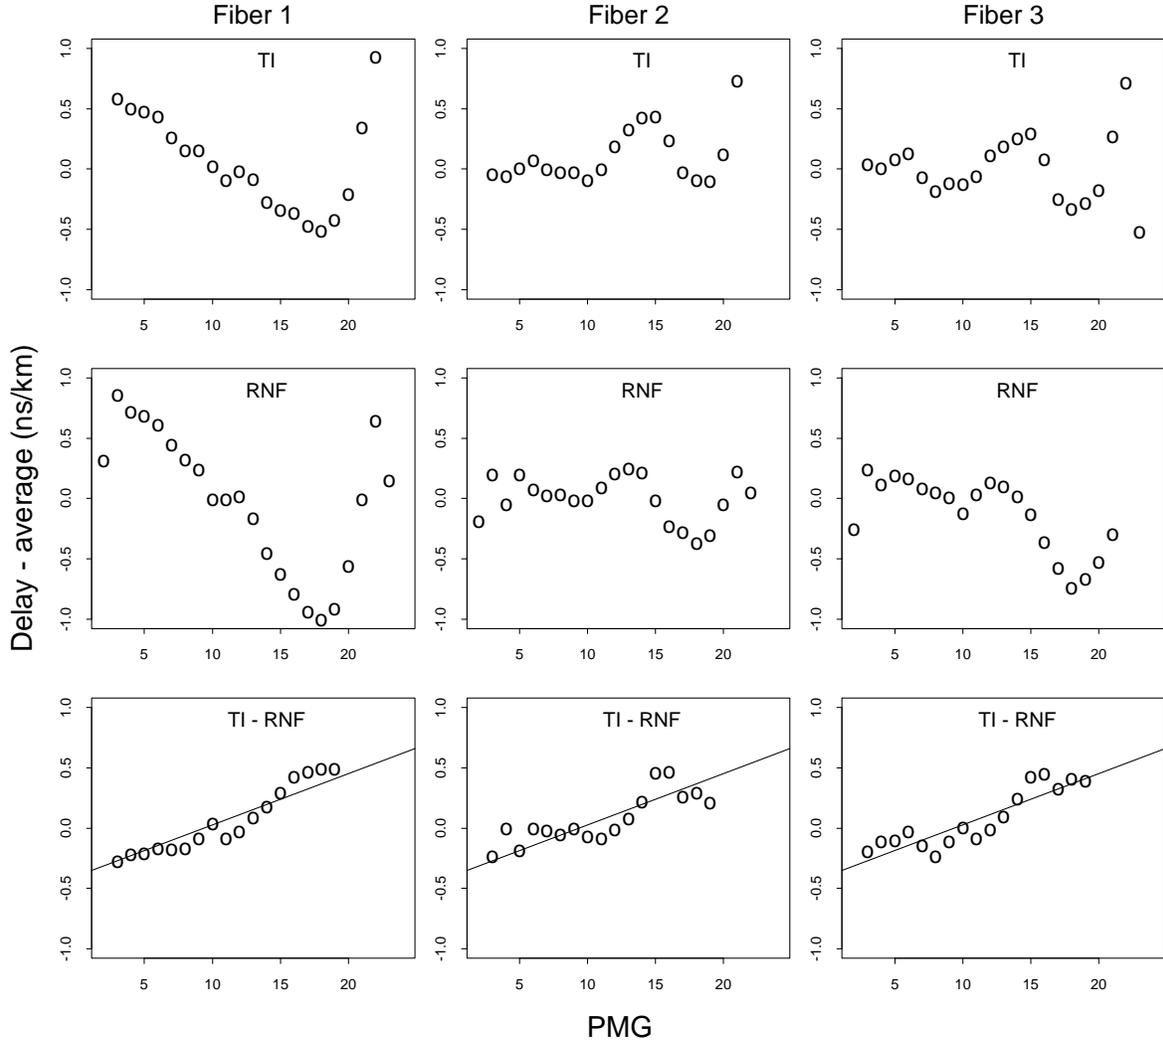


Figure 10: Predicted modal delays of three fibers. The top row is based on pre-processed interferometric (TI) measurements, the second row on averaged near field (RNF) measurements, and the third row is the difference of the top two rows, restricted to those modes believed to propagate (see section 4.4). The line in each panel of the third row is a least squares fit to the data in all three panels of the third row.

the capacity of empirical adjustments so as to avoid overfitting. The result in Figure 10 suggests part of our approach for the problem of adjusting the calculated mode delays in order to predict bandwidth. All uncertainties in the measurements and material properties are incorporated into the single empirical relationship for “corrected delays”,

$$\hat{\tau}_i = \tau_i - i\beta \tag{6}$$

where τ_i is the calculated delay of mode i , β is to be estimated, and $\hat{\tau}_i$ is the corrected delay. An intercept term is not necessary because only *differences* in delays enter into the bandwidth calculation (see equation (7)). There are two reasons for this choice of corrective model. First, as just noted, some smooth correction has to be applied to the delays and a linear adjustment is the simplest choice which fits our data. Secondly, this type of model can approximate the effects of small changes in choices of the material properties of the glass, as reflected through the Sellmeier coefficients, and thus permits compensating for the fact that the true Sellmeier choices are unknown. This observation results from empirical comparisons of the effects of different Sellmeier choices on the mode delay calculations.

Two other effects important to model empirically are differential modal attenuation (DMA) and mode coupling. Our data is poorly suited to estimating DMA, and we did not have access to explicit DMA measurements. DMA measurements in the literature (Olshansky and Oaks, 1978) suggest that we model DMA by a step function: all modes with PMG index below some threshold are assumed to be unattenuated, and power in all other modes is assumed to be lost. The threshold may depend on many factors, but theory suggest that it does not depend strongly on the profile. Because all of the fibers in our study are similar except for their profiles, a single threshold should apply to all of them. We choose the threshold by the following argument. Even in a high bandwidth fiber, the top few PMGs typically have delays that differ significantly from the rest of the PMGs. Such a situation is exhibited in the first column of Figure 10, where PMGs larger than 19 have delays that deviate sharply from the relatively straight line on which the remaining modal delays lie. This is due to the fact that the fields of those modes overlap the cladding significantly. High measured bandwidths are incompatible with these top few modes propagating. Therefore, we choose the threshold to be just below the lowest mode that is significantly influenced by the cladding.

<i>Cross-validation</i>	<i>Correlation coefficient</i>
None	0.63
Random 7-fold	0.61
Leave one preform out	0.56

Table 2: Performance of the fiber bandwidth model under various cross-validation procedures.

Propagation may also be influenced by mode coupling, in two ways. The first is a loss of power to radiation among modes with high PMG index, causing them to have higher attenuation. For our purposes, this effect is folded into the above DMA model. The other effect is a “blurring” of the impulse response, which can improve bandwidth. We elect to ignore this second effect, as it is generally believed to be small in current fibers.

Summarizing, our method for relating a measured profile to a predicted bandwidth is to pre-process the profile, calculate the modal delays, adjust these delays with an empirical correction, ignore the high order modes, and finally deduce the bandwidth from the delays of the low order modes. This final step is accomplished from the formula

$$\text{BW (MHz km)} = \frac{170}{\sqrt{\sum_{i=1}^M (\hat{\tau}_i - \bar{\tau})^2}} \quad (7)$$

where M is the threshold PMG index, $\bar{\tau}$ is the mean of $\hat{\tau}_1, \dots, \hat{\tau}_M$, and all delays have units of ns/km. The formula (7) for predicting bandwidth from RMS pulse width has been shown empirically to be a very good approximation to the usual definition of bandwidth given in Section 3.1. The model has one free parameter, β from equation (6), which we fit to measured bandwidths by nonlinear least squares.

To assess this model, we performed a cross-validation procedure similar to that of Section 3.4. We held out the fibers drawn from one of the seven preforms, fit this model to the fibers drawn from the remaining six preforms, and used the results to predict the bandwidths of the fibers that were held out. This procedure was repeated holding out each preform. In addition, we performed seven-fold cross validation where, in each step, a random seventh of the fibers were held out. The results are displayed in Table 2.

The difference between the random seven-fold and the leave-one-preform-out correlation coefficients indicates that there is some information in the preform identification, whereas one

would hope this would not be the case. Recall our premise that the refractive index profile should contain *all* of the information necessary to explain bandwidth differences between fibers, and that given the profiles the preform (or lathe) identification should not help. This situation is similar to the lathe effect discussed in Section 3.4, although smaller here. There are a number of possible explanations; most likely, the measurement bias changed over time, since the order in which the fibers were measured was not randomized.

5 Feedback to Engineering Problems

Our data analysis and modeling began with the goal of solving the engineering problem of predicting bandwidth from preform refractive index profiles. If sufficiently accurate predictions could be made, manufacturing costs could be reduced by making suitable adjustments early in the production process. The empirical modeling work described in Section 3 shows some success, but the cross-validation assessment described in Section 3.4 also shows some remaining lathe dependency and ultimately inadequate predictability. These empirical models demonstrate promise and represent a big step beyond initial expectations of some who were familiar with the manufacturing process and measurement quality, but the results were not adequate to use for tuning the production process or for sorting out bad quality preforms.

To understand and overcome the deficiencies, we moved to several related efforts to construct models on different types of data. Our goals were to establish a “proof of concept” that adequate models of bandwidth from preform profiles should be achievable in the right circumstances and to determine to what extent inaccuracies in preform profiles were limiting predictive ability. To use the physical theory more directly, along with possibly higher quality and more detailed propagation measurements, we focused on modeling mode delays based on *fiber* refractive index profiles as described in Sections 4.3 and 4.4. These models provide convincing “proof of concept” and validate our modeling strategy incorporating an interplay between sophisticated empirical and physical-theoretic approaches at each step.

Unfortunately, the model for predicting *fiber* mode delays from fiber refractive index profiles does not add direct engineering value. But this modeling success based on fiber profiles provides strong evidence that similar success should also be obtainable using *preform* profiles.

Our modeling results convinced those involved that some poorly-understood problems led to degraded accuracy in measurements of the preform profiles. Extra resources were therefore directed into a careful investigation to assess and improve many aspects of the preform measurement system.

After upgrading the profiling equipment, refractive index profiles were taken from several positions and angles on a section of a preform which was then drawn into 4 km of fiber. The ends of the resulting fiber were also profiled along several diameters using the refractive near field technique. The averaged preform and fiber profiles were used to predict the mode group delays of the fiber in a manner similar to that described in Section 4. This time, however, *no preprocessing* was done to the fiber and preform profiles beyond averaging over angles and smoothing. Also, *no correction* was made to the predicted mode delays, as was done by estimating β in equation (6).

By this point the engineering team was not so much interested in predicting bandwidth, a scalar measurement, as it was in predicting the functional dependence of propagation delay verses principal mode group index. The PMG delays are more informative and fundamental than bandwidth. A comparison between the measurements and the predictions from the two types of profiles is presented in Figure 11. The agreement for PMG $\lesssim 25$ is remarkable. The data that yielded the measured mode group delays do not contain information on the PMGs greater than 25. A large discrepancy exists in the lowest two PMGs and slight differences in the trend at higher PMGs. The former difference is a failure of the predictions, as the data are unambiguous near the center of the fiber, while the origin of the latter is less certain. Overall, however, the level of agreement (< 0.5 nm/km) is easily high enough to be of engineering significance.

Thus, our modeling approach led—albeit indirectly and through several modeling steps—to demonstrating in Figure 11 that measured preform profiles can be used to predict a basic feature of fiber transmission properties, namely mode delays.

Both empirical and physical-theoretic modeling approaches provided important steps towards this result. The empirical approach had the potential to deliver quick improvements to manufacturing bandwidth yields, and it succeeded in convincing doubters that preform profile measurements can indeed help to predict fiber bandwidths, even if prediction accuracy was still an issue. The physical approach, on the other hand, served more naturally to

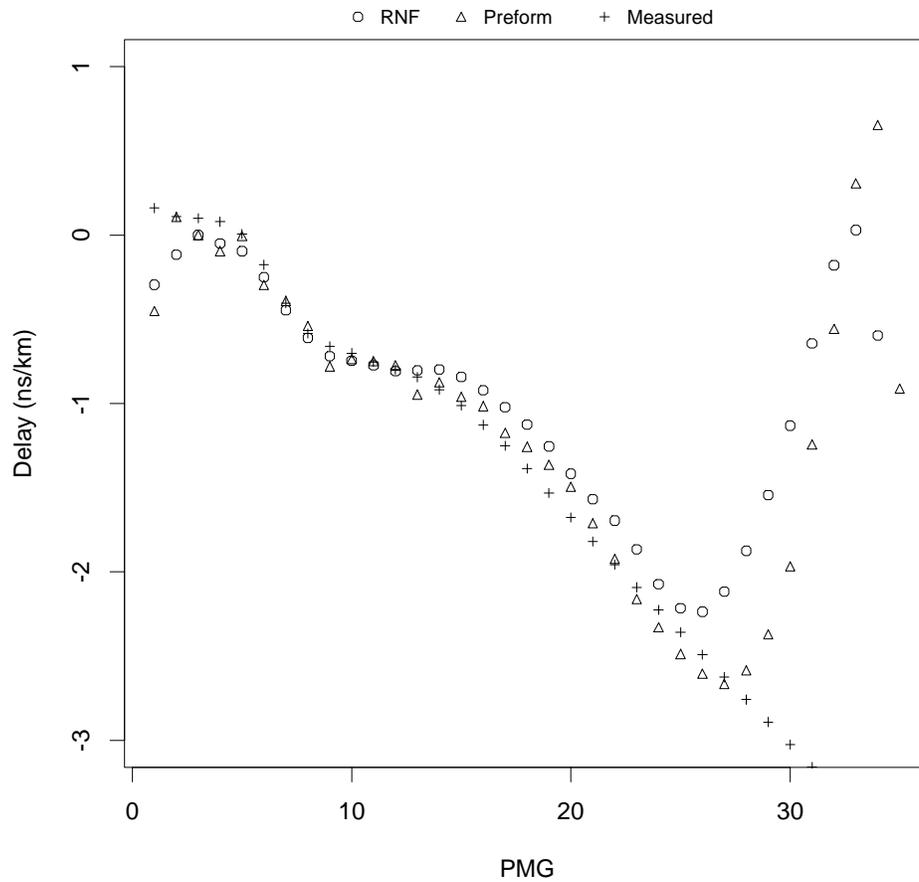


Figure 11: Mode group delays as predicted from RNF fiber profiles (RNF), preform profiles (Preform), and measured by propagating pulses in the various mode groups through a 4 km length of fiber (Measured). The data do not contain information on $\text{PMG} \gtrsim 25$.

indicate where further work was needed to improve prediction accuracy because it produces ground-truth outcomes: a given refractive index profile will result in a computable set of mode delays subject to certain assumptions and approximations that can be listed. In particular, the physically-based model suggested that preform profile measurements were grossly inaccurate and this was confirmed by the close agreement obtained in Figure 11 once the profiling equipment was upgraded. A fundamental message that emerged from both modeling approaches is that each benefits from a measure of the other. Whichever approach is primary—empirical or physical—it remains important to blend physical theory with statistical empirical techniques.

A further byproduct of this work has been to validate the use of multimode propagation models in the design of future systems. Figure 11 is a convincing demonstration of the adequacy of the physical assumptions and approximations, and is therefore convincing proof that simulations of optical transmission systems that employ the physical models can be trusted. Such simulations were influential in the design of 10 Gb systems employing multimode fiber, and also were instrumental in the development of laser and multimode fiber specifications that were included in the IEEE 10 Gigabit Ethernet standard (Golowich, Kolesar, Ritger and Pepeljugoski, 2001).

Acknowledgments

The authors are grateful to W. A. Reed and A. J. Ritger for supplying the data studied in this paper and for numerous useful discussions.

References

- Becker, R. A., Chambers, J. M. and Wilks, A. R. (1988). *The New S Language*, Wadsworth Publishing Co, Belmont CA.
- Buckler, M. (1982). Measurement of bandwidth versus impulse response width in multimode fibers, *Technical Digest – Symposium on Optical Fiber Measurements*, Vol. NBS-SP-641, pp. 33–36.
- Fleming, J. W. (1978). Material dispersion in lightguide glasses, *Elect. Lett.* **14**(11): 326–328.

- Golowich, S. E., Kolesar, P. F., Ritger, A. J. and Pepeljugoski, P. (2001). Modeling and simulations for 10 Gb multimode optical fiber link component specifications, *Proceedings of OFC*, Anaheim, CA.
- Kiiveri, H. T. (1992). Canonical variate analysis of high-dimensional spectral data, *Technometrics* **34**: 321–331.
- Lenahan, T. A. (1983). Calculation of modes in an optical fiber using the finite element method and eispack, *The Bell System Technical Journal* **62**(9): 2663–2694.
- Marcuse, D. (1981). *Principles of Optical Fiber Measurements*, Academic Press, New York.
- Olshansky, R. (1979). Propagation in glass optical waveguides, *Reviews of Modern Physics* **51**(2): 341–367.
- Olshansky, R. and Oaks, S. M. (1978). Differential mode attenuation measurements in graded-index fibers, *Appl. Opt.* **17**(11): 1830–1835.
- Ramsay, J. and Silverman, B. (1997). *Functional Data Analysis*, Springer-Verlag, New York.
- Schumaker, L. L. (1981). *Spline Functions: Basic Theory*, Krieger, Malabar, Florida.
- Snyder, A. W. and Love, J. D. (1983). *Optical Waveguide Theory*, Chapman and Hall, London.