PUBLICATIONS OF THE ASTRONOMICAL SOCIETY OF THE PACIFIC, **115**:218–234, 2003 February

Removing the Fringes from Space Telescope Imaging Spectrograph Slitless Spectra[1]

ELIOT M. MALUMUTH,[2,3,4] ROBERT S. HILL,[2,3,4] TED GULL,[2] BRUCE E. WOODGATE,[2] CHARLES W. BOWERS,[2]
RANDY A. KIMBLE,[2] DON LINDLER,[2,5,6] PHIL PLAIT,[2,5,7] AND MORLEY BLOUKE[8]
Received 2002 July 31; accepted 2002 October 17; published 2003 January 6

ABSTRACT. Using what is known about the physical and chemical structure of the CCD detector on the Space Telescope Imaging Spectrograph (STIS) and over 50 calibration images taken with different wavelength mappings onto the detector, we have devised a model function that allows us to predict the fringing of any spectral image taken with the STIS CCD. This function is especially useful for spectra taken without a slit with the G750L grating. The STIS parallel observing program uses this "slitless spectroscopy" mode extensively. The arbitrary mapping of wavelength versus position that results from each source's chance position in the field renders direct calibration of the fringe amplitudes in this mode impossible. However, we find that correcting observed data using our semiempirical fringing model produces a substantial reduction in the fringe amplitudes. Tests using the flux calibration white dwarf standard G191-B2B show that we can reduce the fringe amplitude in the 9000–10000 Å region from about 20% peak to peak (10% rms) to about 4% peak to peak (2% rms) using the model, while a standard calibration using a "fringe flat" reduces the fringe amplitudes to 3.3% peak to peak (1.7% rms). The same technique is applicable to other astronomical CCDs.

1. INTRODUCTION

Fringing is an annoying fact of life for many astronomical CCD detectors. The CCD employed in the *Hubble Space Telescope*'s (*HST*'s) Space Telescope Imaging Spectrograph (STIS; a 1024 × 1024 pixel, backside-thinned, UV-enhanced device developed by Scientific Imaging Technologies; Kimble et al. 1994) is no exception. Fringing is a rapidly varying sensitivity as a function of wavelength resulting from interference among the incident and reflected beams within the CCD layers; it thus generally grows in amplitude toward the longer wavelengths of CCD use, as the declining absorption coefficient of silicon makes a greater proportion of the light available for such interference instead of being absorbed on the first passage through the detection layer. In STIS, dispersed spectra show the onset of fringing in the neighborhood of 7000 Å, with peak-to-peak amplitudes of the sensitivity modulation reaching ∼25% at 9800 Å (Kimble et. al. 1998). The fringing of STIS spectra is particularly severe as a result of the near collimation (f/36) of the beam incident on the CCD (fringing effects are decreased in faster optical systems).

For spectra taken with a slit defining the input position, the resulting sensitivity can be calibrated and removed in a straightforward manner. A "fringe flat" is obtained by observing a tungsten continuum lamp through a narrow slit and dividing out the broad spectral shape of the lamp from the resulting image.Figure 1 shows a fringe-flat image and a plot of the sum of the central three rows taken from a typical G750L fringe flat. The nominal observation strategy for STIS when using the G750L grating is to obtain a fringe flat at the time of the observations. However, spectra can be recorded by STIS over an extended field with no defined aperture. This mode is known as "slitless spectroscopy" and is a prime observing mode for STIS parallel observations. All spectra obtained in this manner will suffer from fringing, but there is no possibility of placing a narrow slit at the position of each source to acquire directly the appropriate fringe flat as described above.

Although we can not acquire fringe calibration data at all possible mappings of the wavelength versus position on the CCD detector, we show in this paper that a semiempirical fringing model developed to fit a modest set of calibration exposures at different wavelength settings can be used to predict and remove the interference fringes from spectra at any wavelength setting. The resulting fringe model, applied to slitless spectra, corrects for the fringing behavior with an accuracy near that achieved by applying the appropriate fringe flat for a spectrum taken through a slit.

[1] Based on observations with the NASA/ESA *Hubble Space Telescope*, obtained at the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-26555. Funding of this activity was through the Space Telescope Imaging Spectrograph Guaranteed Time Observations.

[2] Laboratory for Astronomy and Solar Physics, Code 681, NASA Goddard Space Flight Center, Greenbelt, MD 20771; Eliot.Malumuth@gsfc.nasa.gov.

[3] Raytheon ITSS Corp., Lanham, MD 20706.

[4] Science Systems and Applications, Inc., Lanham, MD 20706.

[5] Advanced Computer Concepts, Potomac, MD 20854.

[6] Sigma Space Corporation, 9801 Greenbelt Road, Lanham, MD 20706.

[7] Department of Physics and Astronomy, Sonoma State University, Rohnert Park, CA 94928.

[8] Scientific Imaging Technologies, Inc., 10500 SW Nimbus Avenue, Tigard, OR 97223-4310.

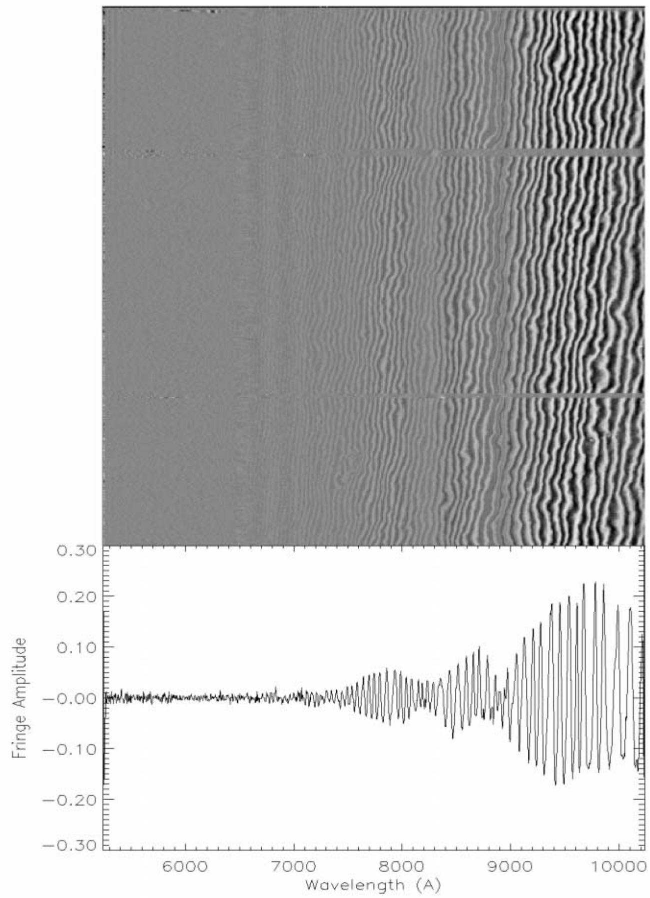

FIG. 1.—Fringe-flat image that has been normalized by a spline fit to the lamp spectral output and stretched from −30% to +30%. The lower panel is a plot of the mean of the three rows 511, 512, and 513 of this 1024 × 1024 pixel image. The fringe pattern has nodes at 6900, 7400, 8200, and 8900 Å. These nodes are vertical in the image, even though the fringes are tilted.

A careful examination of the image in Figure 1 shows several interesting and pertinent facts. One, although the fringes are basically vertical, they do tilt a little so that they are farther to the right near the top of the image than at the bottom. This indicates that the thickness of the CCD's detection layer is wedge shaped. Two, the fringes have kinks and wiggles in them, indicating that there are local variations in the CCD's detection layer thickness. And three, the "washed out" areas (nodes of small fringe amplitude) are vertical. As we will see in § 3, this indicates that the lower layers of the CCD are close to constant thickness. These observations suggest that we need to solve for the thickness of the detection layer on a pixel-to-pixel basis, while we may use a constant thickness for the other layers in the CCD.

On 1998 July 29, 40 fringe flat/wavecal pairs were obtained at nonstandard mode select mechanism (MSM) positions (i.e., varying the grating tilts to produce different central wavelengths). These data, in addition to 11 fringe flat/wavecal pairs

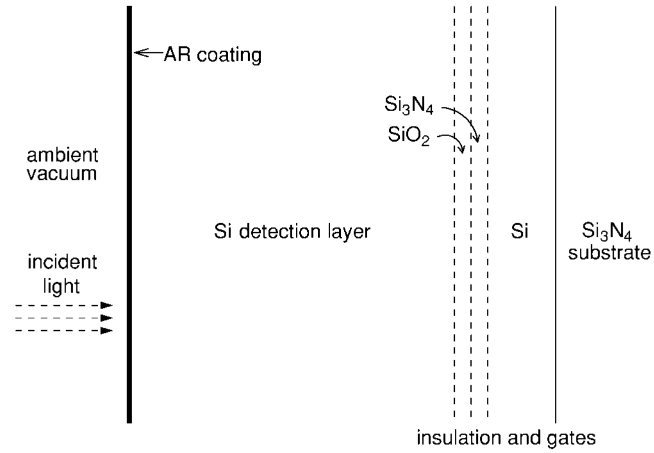

FIG. 2.—Schematic depiction of STIS CCD structure. Relative thicknesses of the layers are not to scale. Although the insulation and gate layers are planar in our model, they have a complex three-dimensional structure in reality (e.g., see Figs. 1.13 and 1.14 of Janesick 2001).

taken at the standard MSM position (but with wavelength shifts due to MSM nonrepeatability and thermal variations of a few pixels), make up the basic data set used to develop, test, and calibrate our model. This technique can be used to make similar models of other astronomical CCDs.

In § 2, we discuss the basic physics of CCD fringing in general and of the STIS CCD in particular. In § 3, we discuss the model of the STIS CCD fringing, and we use the observed data to quantify the parameters of the semiempirical model. In § 4, we test how well the model works on STIS calibration data and apply the model to real STIS parallel slitless spectra.

2. STIS FRINGING IN GENERAL

2.1. CCD Structure

CCD fringing is an instance of multilayer thin-film interference, well known in applied optics. A schematic depiction of the structure of the STIS CCD is shown in Figure 2. Shortward of 7000 Å, the doped-silicon detection layer absorbs photons efficiently, and most of the photons are absorbed near the surface of the chip. However, longward of 7000 Å, the coefficient of absorption of silicon falls rapidly with increasing wavelength. Thus, significant numbers of long-wavelength photons traverse the detection layer completely and are reflected back and forth between its two boundaries. The crossing and recrossing waves mutually interfere. The dependence of phase on wavelength produces the fringe pattern seen in Figure 1.

An important feature of the fringe pattern is caused by the three thin layers labeled as the insulation and gate structure in Figure 2. Omission of these layers would result in the simple fringe pattern shown in Figure 3. Here, the strength of the fringes is determined primarily by the coefficient of absorption of silicon. However, the fringe pattern in Figure 1 has distinct lobes in addition to the overall taper. These lobes are an additional interference pattern arising in the three thin layers. This interference affects the waves that are reflected back into the detection layer.

For the purpose of this paper, perhaps the most important feature of the observed fringe pattern is its irregularity. In Figure 1, many apparently random variations are seen in the intensity and spacing of adjacent fringes. This effect signifies pixel-to-pixel variations in the physical CCD structure. Therefore, in order to completely characterize the STIS CCD, a map of the chip must be produced by fitting observed fringes with multilayer optical models on a pixel-by-pixel basis. The computational method is described below, and the details of parameter fitting are explained in § 3.

2.2. Fresnel Equation Model

A convenient computational method for solving thin-film problems is offered by the Fresnel equations, which are derived by enforcing continuity of the electromagnetic field components across the boundaries between the layers.

Our version of the method is adapted from that embodied in IMD, a visualization tool written in Interactive Data Language (IDL) by Windt (1998). The inputs to the model are the parameters of each layer in the stack; the most important ones are the thickness and the complex refractive index $\boldsymbol{n}$ as a function of wavelength. The complex refractive index $\boldsymbol{n}$ is defined as $n + ik$, where n is the refractive index and k is the extinction coefficient. Although this model is precise in the sense that the whole infinite series of interfering waves is included implicitly, it still entails approximations and simplifications, as described below.

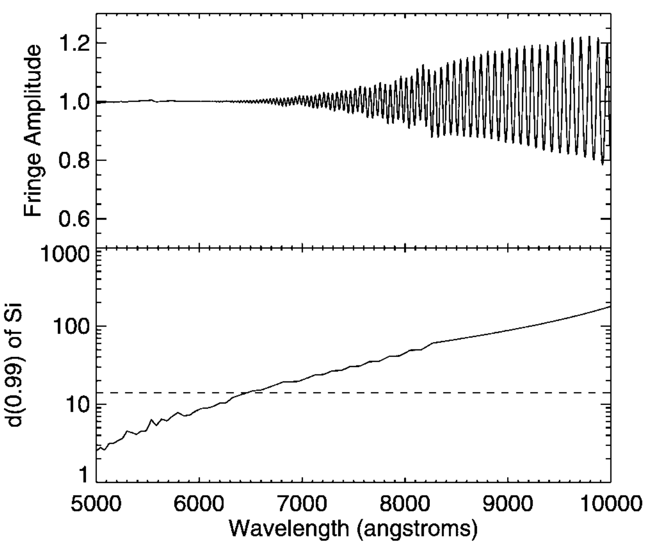

FIG. 3.—*Top:* Normalized fringe pattern for a single silicon layer of thickness 14 μm in ambient vacuum on a substrate of Si_3N_4. The envelope of the pattern is determined by the coefficient of absorption of Si, which is low at wavelengths longer than ~7000 Å. The lobed structure (slowly varying modulation of fringe envelope) in the real CCD fringing is absent because this model omits the insulation and gate layers, which introduce phase differences into the light reflected from the back of the chip. *Bottom:* Depth in μm at which 99% of the incident light is absorbed in a thick block of silicon. The dashed line shows a depth of 14 μm, i.e., the nominal thickness of the STIS CCD and the thickness used to compute the fringes shown in the top panel. Fringing is apparent at wavelengths for which as little as 1% of the light reaches the back of the CCD.

TABLE 1
STRUCTURE OF STIS CCD

Layer (n) (1)	Purpose (2)	Material (3)	Nominal Thickness[a] (d_n) (μm) (4)	Model Thickness (d_n) (μm) (5)	Roughness (μm) (6)
0	Ambient medium	Vacuum	Infinite	Infinite	...
1	AR coating	MgF_2	0.06	0.06	0.01
2	AR coating	Si_3N_4	0.03	0.03	0.02
3	Detection	Si	13–15	13–15	0.06
4	Insulation	SiO_2	0.07	0.28	0.06
5	Gate	Si_3N_4	0.06	0.24	0.08
6	Structures	Si	0.40	0.81	0.10
7	Substrate	Si_3N_4	Effectively infinite	Infinite	0.11

[a] Nominal thickness provided by SITe.

The structure of the algorithm relies on the existence of a preferred direction, which is the direction of the incident light. The *top* layer is defined as the first layer encountered by an incoming photon, the *ambient medium* is that in which the incoming photon travels before encountering the CCD, and the *substrate* is an infinitely thick medium beneath the bottom layer of the stack. In the following formulae, N is the number of layers in the stack. The number of interfaces is therefore $N + 1$, since both the interface between the top layer and the ambient medium and that between the bottom layer and the substrate are included in the model.

The algorithm builds up the stack one layer at a time, beginning with the bottom layer. Formally, each step of the algorithm is as follows:

$$r_k = r[r_{k+1},\ r(\boldsymbol{n}_k,\ \boldsymbol{n}_{k+1}),\ \boldsymbol{d}_{k+1}\boldsymbol{n}_{k+1}],$$

$$\boldsymbol{t}_k = \boldsymbol{t}[\boldsymbol{t}_{k+1},\ \boldsymbol{t}(\boldsymbol{n}_k,\ \boldsymbol{n}_{k+1}),\ \boldsymbol{d}_{k+1}\boldsymbol{n}_{k+1}],$$

where

r_k is the resultant amplitude reflection coefficient, taking into account layers k through N;

$r(\boldsymbol{n}_k,\ \boldsymbol{n}_{k+1})$ is the amplitude reflection coefficient for the interface between layers k and $k + 1$;

t_k is the resultant amplitude transmission coefficient, taking into account layers k through N;

$t(\boldsymbol{n}_k,\ \boldsymbol{n}_{k+1})$ is the amplitude transmission coefficient for the interface between layers k and $k + 1$;

$\boldsymbol{n}_{k+1}$ is the complex refractive index of layer $k + 1$; and

d_{k+1} is the thickness of layer $k+1$.

All of $\boldsymbol{n}$, r, and t are complex and a function of wavelength; i.e., they include phase effects as well as amplitude effects in the ordinary sense.

The iteration does not require any special terminating step when the ambient medium is reached. In essence, each layer i, as it is added to the stack, is an ambient medium in relation to the layers that already exist. The thickness of layer $i+1$ becomes determinate only as layer i is added to the top. When $i = 0$, the true ambient medium is included and the algorithm is finished, so in our notation r_0 and t_0 are the amplitude reflection and transmission coefficients, respectively, for the whole stack.

The iteration is initialized using r and t for the bottom interface:

$$r_N = r(\boldsymbol{n}_N, \boldsymbol{n}_{\text{substrate}}),$$

$$\boldsymbol{t}_N = \boldsymbol{t}(\boldsymbol{n}_N, \boldsymbol{n}_{\text{substrate}}).$$

The actual formulae embodied in our software are as follows:

$$r_k = \frac{r(\boldsymbol{n}_k, \boldsymbol{n}_{k+1}) + \boldsymbol{r}_{k+1} \exp(\boldsymbol{2i}\beta_{k+1})}{\boldsymbol{1} + \boldsymbol{r}(\boldsymbol{n}_k, \boldsymbol{n}_{k+1})\boldsymbol{r}_{k+1} \exp(\boldsymbol{2i}\beta_{k+1})}, \tag{1}$$

$$\boldsymbol{t}_k = \frac{\boldsymbol{t}(\boldsymbol{n}_k, \boldsymbol{n}_{k+1})\boldsymbol{t}_{k+1} \exp(\boldsymbol{i}\beta_{k+1})}{\boldsymbol{1} + \boldsymbol{r}(\boldsymbol{n}_k, \boldsymbol{n}_{k+1})\boldsymbol{r}_{k+1} \exp(\boldsymbol{2i}\beta_{k+1})}, \tag{2}$$

where $\beta_{k+1} = 2\pi d_{k+1}\boldsymbol{n}_{k+1}$ for normal incidence, which is assumed throughout this paper.

At some cost in awkwardness, this notation is deliberately more explicit than that of Windt (1998), in order to make clear how the addition of each layer works. The relationship between the layers and the interfaces is shown by equations (57) and (58) in chapter 1 of Born & Wolf (1999), which give the simple case of one layer between two semi-infinite media. The results of our implementation have been checked against those of Windt's IMD program for many specific cases.

To compare with data, the results need to be in terms of power rather than amplitude:

$$T = \text{Real}\left\{\frac{\boldsymbol{n}_{\text{substrate}}}{\boldsymbol{n}_{\text{ambient}}}\right\} |\boldsymbol{t}_0|^2,$$

$$\boldsymbol{R} = |\boldsymbol{r}_0|^2,$$

$$\boldsymbol{A} = \boldsymbol{1} - \boldsymbol{T} - \boldsymbol{R},$$

where T is the power transmission coefficient for the whole stack, R is the power reflection coefficient for the whole stack, and A is the power absorption coefficient for the whole stack.

Windt (1998) describes several possible extensions of the Fresnel equation method. One is to differentiate forward-going and backward-going waves and to compute the electric field intensity on a grid of depths in the chip. This refinement permits computation of the power absorbed by the detection layer alone. However, the resulting algorithm is too time consuming to be used in fitting 10^6 pixels, so instead we use A for the stack. This leads to errors in the detected power of up to a few percent at wavelengths of $\sim$0.9–1 μm, but this error is negligible compared to the relative fringe contrast that we fit.

Another refinement of the thin-film model is to include roughness of the interfaces or a zone of blending between each pair of layers. Windt (1998) describes several parametric formulae for such transition zones. In our software, we implement only one of these possibilities, namely, the transition profile in the form of the error function (erf). The characteristic thicknesses of these transition zones at every interface are fixed parameters of our model for the STIS CCD. The effect of a nonideal interface is loss of specular reflectance (Windt 1998), reducing fringe contrast. In reality, the light scattered by nonideal interfaces is absorbed elsewhere in the chip, so treating the scattering as pure loss is again a simplification.

The formulae for amplitude reflection and transmission at each interface are substituted into the algorithm described above. For normal incidence,

$$r(\boldsymbol{n}_k, \boldsymbol{n}_{k+1}) = \tilde{\boldsymbol{w}} \frac{\boldsymbol{n}_k - \boldsymbol{n}_{k+1}}{\boldsymbol{n}_k + \boldsymbol{n}_{k+1}},$$

$$\boldsymbol{t}(\boldsymbol{n}_k, \boldsymbol{n}_{k+1}) = \frac{\boldsymbol{2n}_k}{\boldsymbol{n}_k + \boldsymbol{n}_{k+1}}.$$

The parameter $\tilde{w}$ describes the decrease in specular reflection due to interfacial roughness. For the error function profile used here (Windt 1998),

$$\tilde{w} = \exp\left[-\frac{\sigma^2}{2}\left(\frac{4\pi}{\lambda}\right)^2\right],$$

where $\tilde{w}$ is the Fourier transform of the derivative of the profile as a function of depth; σ is the "roughness parameter," i.e., the scale length of the interface zone; and λ is the wavelength of the incident light.

3. FITTING THE THICKNESSES OF THE STIS CCD LAYERS

The preparation for running the fit for the STIS CCD was first to search parameter space by trial and error in the neighborhood of the nominal layer thicknesses (given in col. [4] of Table 1). A fringing pattern was sought as close as possible to that observed in a typical flat field. In particular, an attempt was made to match two features of the fringe spectrum: (1) the period of the high-frequency fringes and (2) the shape

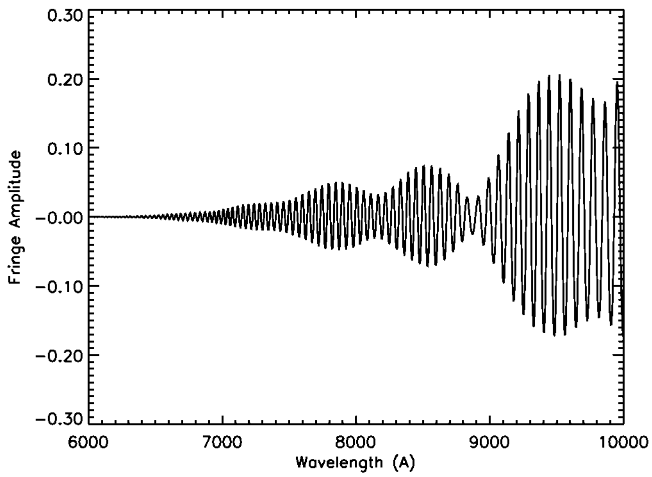

FIG. 4.—Fringe amplitude of the full model for the central pixel using the derived values for each layer's thickness and roughness. These values are not unique; different sets of values can give very similar patterns. The quality of the calibration data on hand is not sufficient to distinguish between models with subtle differences. The pattern is much like the observed fringe flats (see Fig. 1), except that the fringe spacing is very regular here and it is not in the observed fringe flats. This is because this model is for a single pixel with a single value of d_3 (the thickness of the detection layer), while the fringe flats are for a row of pixels of slowly varying thicknesses within the structure.

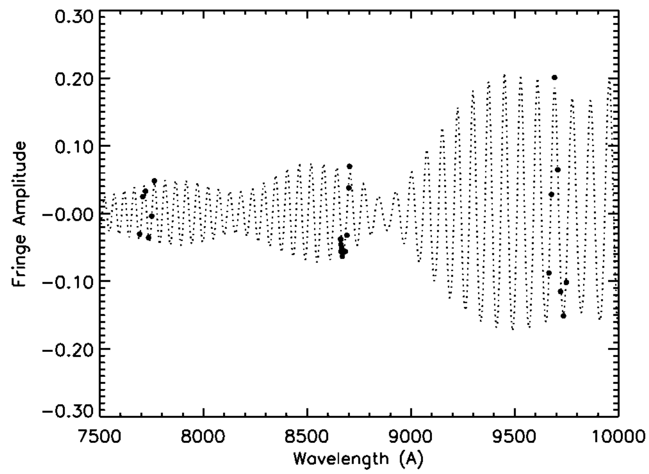

FIG. 6.—Fringe amplitudes between 7500 and 10000 Å for pixel $X = 700$, $Y = 500$ are plotted as large dots. The dotted line is a model for this pixel.

of the overall envelope of the fringes. The period is determined by the thickness of the silicon detection layer, and the shape is determined by the thicknesses of the lower three layers lying between the detection layer and the substrate.

The results of this search are shown in column (5) of Table 1. The primary difference from the nominal parameters is in the lower three layers, which consist of the gates and their insulation from the detection layer. The thin-film model requires these layers to have thicknesses much greater than the nominal values, and it requires interfacial roughness. The discrepancy in thickness may be due to the intricate structure of this part of the CCD (Janesick 2001). Any treatment of the gates and insulation as a plane multilayer is clearly an approximation. However, it allows us to fit the parameter of real interest, namely, the thickness of the detection layer.

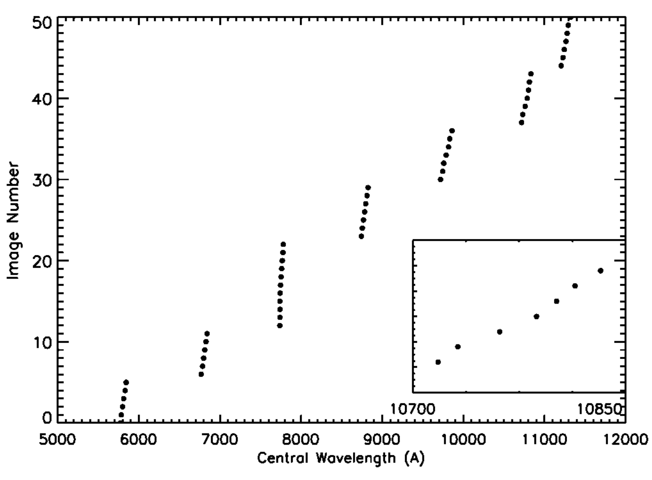

FIG. 5.—Central wavelength (at $X = 512$, $Y = 512$) is shown for the 51 fringe-flat images that compose the calibration data set used in this paper. The 11 images centered near 7800 Å were taken as part of normal calibration programs; the other 40 images were taken at nonstandard MSM positions specifically for this calibration activity. The inset at the lower right is a blowup of the wavelengths near 10750 Å to show the fine sampling.

Figure 4 shows a single-pixel model of the fringing as a function of wavelength using the "model parameters" listed in Table 1, with d_3 set to the fitted value for the central pixel, 14.0206 μm. The model fringe pattern resembles the true fringe pattern (e.g., Fig. 1), with the proper number and wavelength of the nodes. Note that this set of parameters is not unique. Other parameter values that trade changes in the thickness and/or roughness of one of the lower layers for changes in another layer can produce a very similar overall envelope.

Although the fringe pattern in Figure 4 is similar to the observed fringe pattern (Fig. 1) in general, it does not match in detail. The model fringe pattern is regular, while the observed fringe pattern is not. Some fringes are broader and others are narrower, because the STIS detection layer (d_3) has a slowly varying thickness across the chip. In order to model the fringe pattern across the chip, the thickness d_3 of the CCD at each of the over 1 million pixels must be determined.

3.1. The Data

The STIS on-orbit calibration program obtains flat-field exposures in the form of dispersed spectra of a long slit illuminated by an internal continuum calibration source. These "fringe flats" vary in central wavelength because of MSM nonrepeatability and thermal shifts of the optical bench. For the

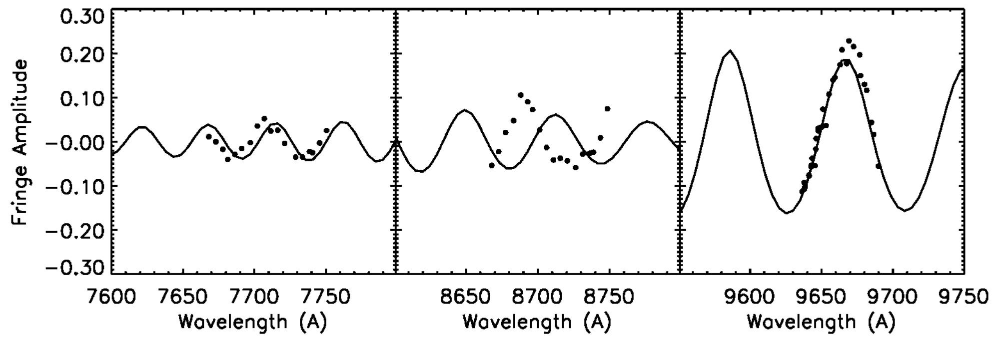

FIG. 7.—Fringe amplitudes for pixels $X = 900$–902, $Y = 499$ are plotted as dots in three separate wavelength regions. The line is from a model with $d_3 = 14.2605\ \mu$m. Although the fit is good at 9670 Å, it is out of phase at 8700 and 7700 Å. No single value of d_3 can be found that is in phase at all three wavelengths given published optical constants for bulk Si.

G750L grating, these few-pixel shifts amount to about $\frac{1}{2}$ fringe (∼40 Å). For small displacements from the nominal slit position, we can adjust the fringe flat. At greater displacements, variations in the CCD structure significantly change the fringe structure, and a simple shift of the fringe flat does not work. Additional fringe flats with significantly different central wavelengths were needed. In 1998 August (proposal 7969), fringe flat/wavecal pairs were obtained at 40 different nonstandard MSM positions (i.e., central wavelength). These data were combined with 11 existing fringe flat/wavecal pairs centered near the "standard" central wavelength to make a fringe-flat cube.

Figure 5 shows the central wavelengths for pixel $X = 512$, $Y = 512$ in the fringe-flat data cube. We selected these central wavelengths to span roughly a whole fringe at seven nearly equally spaced wavelengths across the fringing wavelength range (∼7000–10200 Å).Figure 6 shows the fringe-flat amplitudes for pixel $X = 700$, $Y = 500$ as a function of wavelength. The dotted line is a model for this pixel with d_3 adjusted to get the best fit. In order to make a model that can be used to defringe any and all slitless spectra, we must fit each pixel in the 1024×1024 array and derive the value of d_3 for each pixel.

3.2. Fitting the Pixels

The full model function has 15 parameters (eight thickness and seven roughness values). This is reduced to 13 since the thickness of the ambient medium and the substrate are defined to be infinite. Most of these parameters influence the size and shape of the fringe envelope. Only the thickness of the detection layer d_3 governs the frequency and shape of the high-frequency fringes. Thus, the simplest approach to modeling the STIS CCD is to assume that all of the variables are known except the pixel-to-pixel value of d_3. A trial and error approach allowed us to find values of all of the other thicknesses and roughnesses that lead to a fringe envelope similar to the observed one (see Figs. 1 and 4). Note that our sampling of calibration data for a single pixel is quite sparse (cf. Figs. 5 and 6). This sparseness makes it impossible to solve for additional parameters for each pixel. We initially assumed that the index of refraction and absorption coefficient of the materials were known precisely as a function of wavelength. This assumption was found to be incorrect (see § 3.3).

Even with the simplification of only one variable, solving for d_3 is not straightforward. An analytical solution for d_3 as a function of λ and amplitude is not easily done. In addition, this analysis includes the critical approximations that the input light on any pixel is collimated and monochromatic. The function returns the amplitude given the wavelength and a value of d_3 (holding the thickness and roughness of all other layers constant, and assuming normal incidence). Moreover, there are several equally good values of d_3 given slightly different order numbers, m. Since $2nd_3 \cos\theta = m\lambda$, then $m \sim 100$. With a nominal thickness of $d_3 \sim 14\ \mu$m and an index of refraction of $n \sim 3.6$ at 9000 Å, the order would be approximately $m = 2nd/\lambda = 112$ for normal incidence. Since $\cos\theta$ varies by only one part in 10^4 over the f/36 input beam, the dispersion of 9.8 Å per 2 pixel resolution element implies that $\Delta\lambda/\lambda$ is one part in 1000 for the G750L grating. The effect of assuming normal incidence is negligible. Also, with order ∼100, d_3 can be fitted equally well for $m + 1$ or $m - 1$. We do not have data over a sufficient number of consecutive fringes to isolate the proper order. However, for the practical purpose of computing a fringe flat for a slitless spectrum, this limitation does not matter.

In order to solve for d_3 for each pixel of the CCD, we start with pixel $X = 1020$, $Y = 512$ and guess a value for d_3. We then compute the amplitude at each wavelength represented for that pixel (and its neighbors on either side in the same row) in the fringe flat in the 7000–10200 Å wavelength range. Comparing the model to the observations, we compute the value of χ^2. We then adjust d_3 by a small amount and compute a new χ^2, repeating until we cover the thickness range of one order, about 0.12 μm. The assumed thickness for this pixel is the value of d_3 that minimizes χ^2. Using this value of d_3 as a starting point, we move to the next pixel in the same row ($X =$

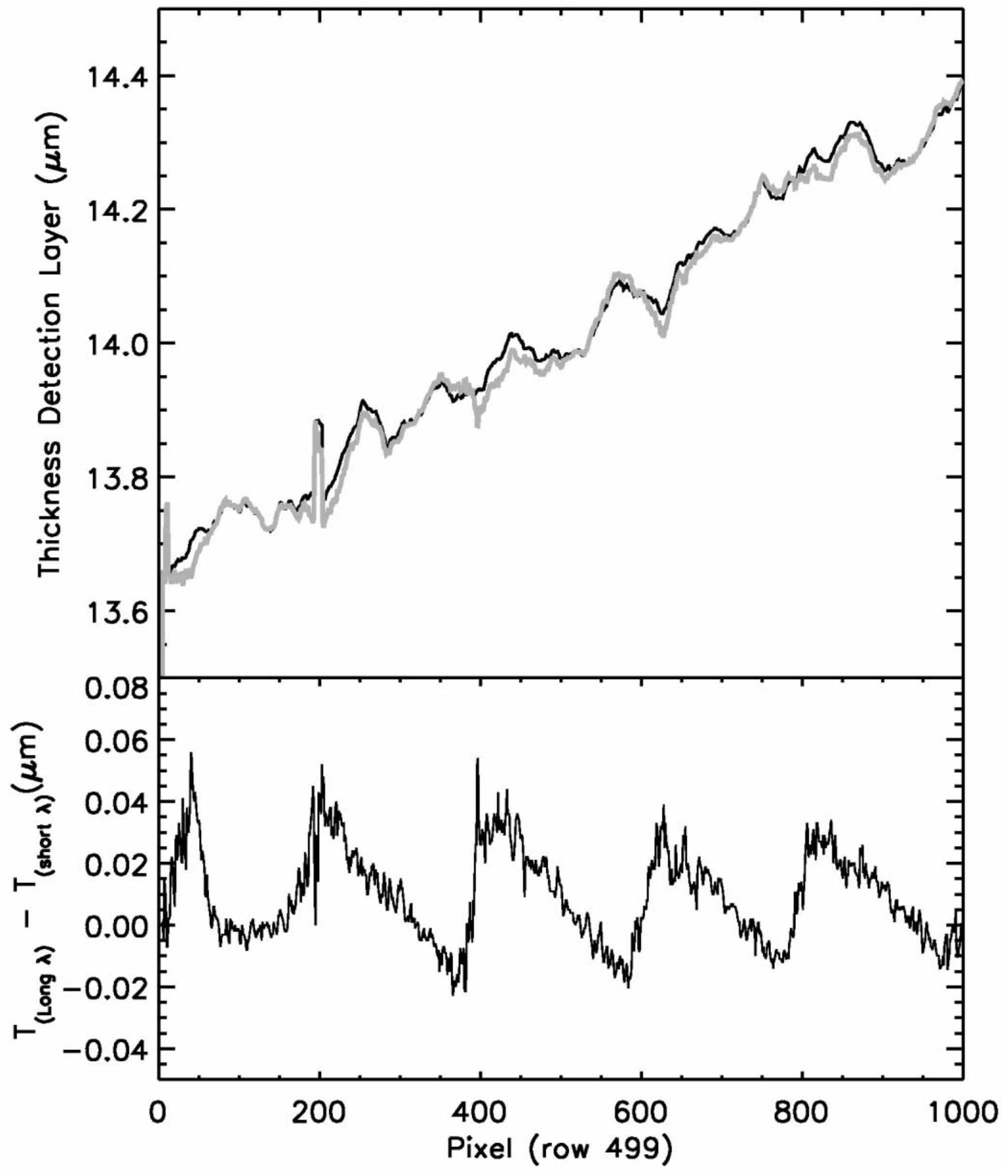

FIG. 8.—Thickness of the STIS CCD detection layer is plotted for row 499. The thin black line is from fitting data between 8900 and 10200 Å only. The thick gray line is from fitting data between 7200 and 8200 Å only. For those pixels where the two lines are identical, there are no data between 7200 and 8200 Å, and the value is set to that of the black line. The lower panel shows the difference between the black line and the gray line. The periodicity with pixel number follows that of the wavelength of the data being fitted, showing that the difference in thickness is a function of wavelength.

1019) and repeat the process, constraining d_3 to be in the same order as it was for pixel 1020. In this way, we work down row $Y = 512$ to pixel $X = 5$. We then move on to the next row ($Y = 513$) and work up the CCD to row 1020. We then return to row 511 and work down the CCD to row 5.

For many pixels, we found that this scheme worked well. However, fits for some pixels yielded no single value of d_3 for all of the wavelengths. This is illustrated in Figure 7, a plot of the data for pixels $X = 900$–902, $Y = 499$ at three wavelengths within the fringing range. A fringe model with $d_3 = 14.26$ μm is shown. The model fits the data quite well at 9670 Å, but it is completely out of phase at 8700 Å, and then it is only slightly out of phase at 7700 Å. No single value of d_3 that fitted well in all three wavelength regimes could be found for many pixels. As shown in the next section, this was due to apparent small discrepancies in the index of refraction of silicon as a function of wavelength. A slight correction to the refractive index at some wavelengths resulted in a good fit for these pixels in all three wavelength regimes.

3.3. The Index of Refraction of Silicon

The phase difference produced by each layer of the stack is dependent on the product of the thickness of each layer times

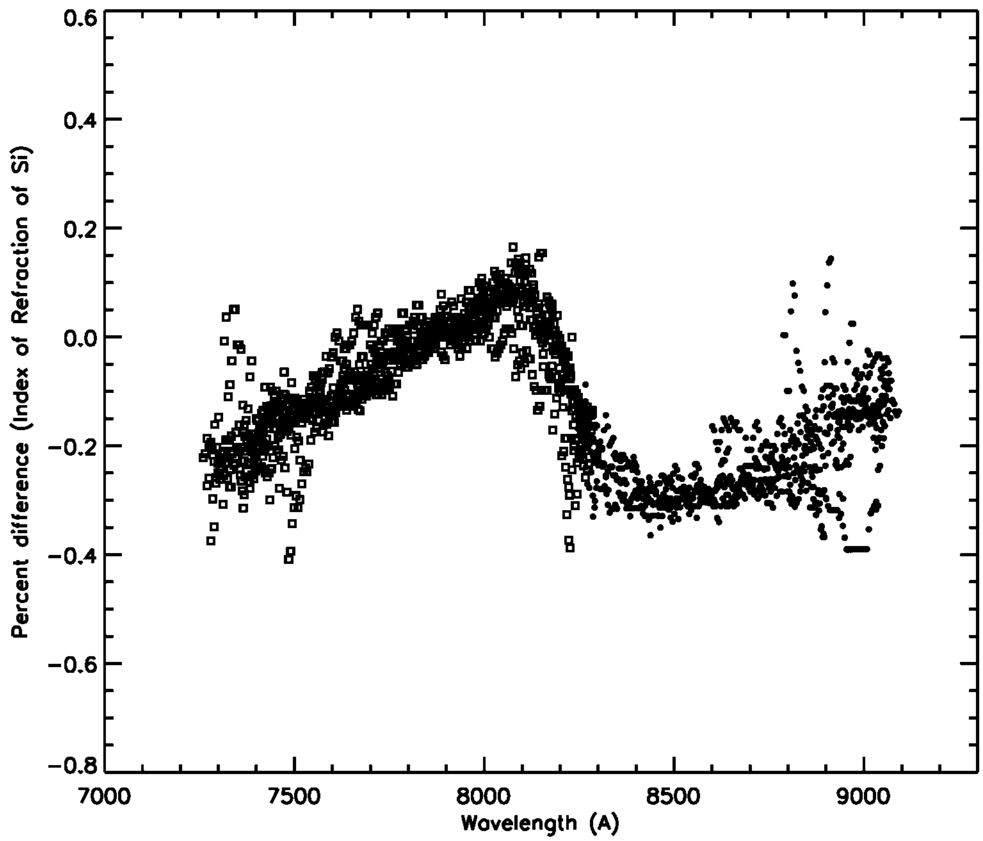

FIG. 9.—Percentage adjustment of the index of refraction of silicon required to have a single value of detection layer thickness (d_3) fit the fringing behavior at all wavelengths for a given pixel. The boxes are derived from the short-wavelength regime for pixels in row 499. The dots are derived from pixels in the middle-wavelength regime for the same row.

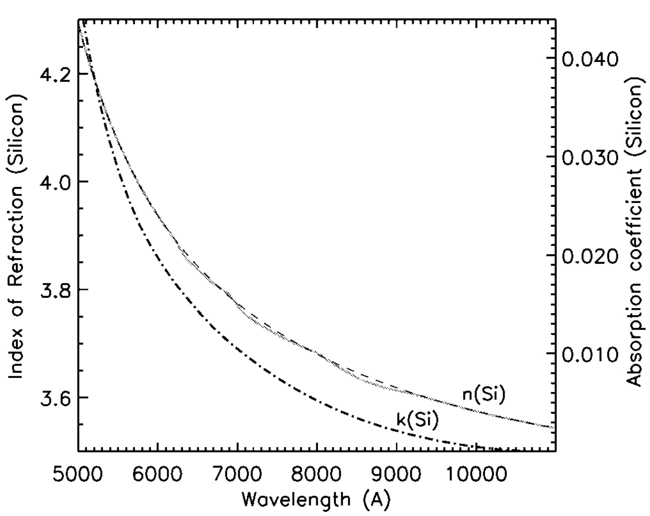

FIG. 10.—Index of refraction of silicon as a function of wavelength. The HOC III values are plotted as the dashed line. Our revisions are shown as the gray line. The revision is normalized to the HOC III values for wavelengths greater the 9000 Å. The absorption coefficient (k) of silicon is shown as the dot-dashed line.

the index of refraction of that layer. The index of refraction is a function of wavelength, for the materials used in the STIS CCD. But how precisely are these values known? In particular, the index of refraction of silicon is of great importance. There are several references for the index of refraction of silicon as a function of wavelength, which have slightly different values. We chose to use the values in the *Handbook of Optical Constants of Solids III* (Palik 1998, hereafter HOC III).

It is possible that the detection layer of the STIS CCD is not pure silicon and that silicon with dopants has a slightly different index of refraction from pure silicon. It is also possible that the index of refraction of silicon is not known with sufficient accuracy at the relevant wavelengths. The supposition that the index of refraction of the STIS CCD is slightly different from that reported for pure silicon in HOC III can be tested by fitting the fringes in the different wavelength regions independently and comparing the derived thicknesses (d_3). The flat-field data frames provide a common wavelength region for all pixels for $\lambda \geq 8900$ Å. We therefore use the derived thickness d_3 for $\lambda \geq 8900$ Å as a baseline.Figure 8 shows the derived thickness of the detection layer in the long wavelength (*black line*) and the short wavelength, $\lambda \leq 8200$ Å (*gray line*), regimes for row 499. For some pixels, there were not enough data points

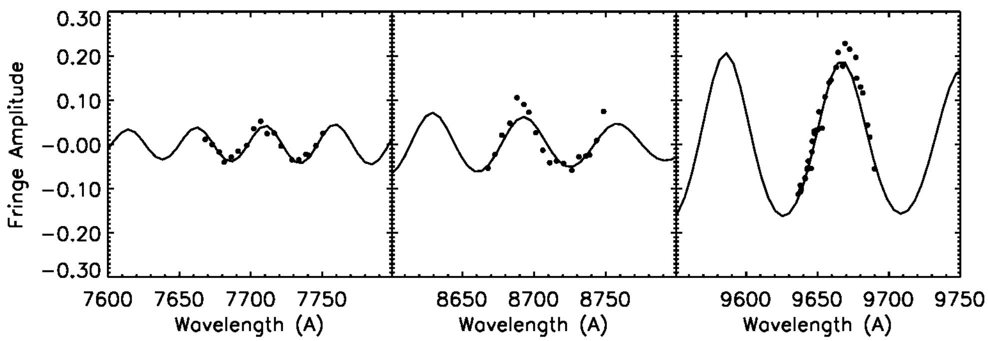

FIG. 11.—Same fringe amplitudes as in Fig. 7, but now compared to a model using the revised values of the index of refraction of silicon. The fit is good at all three wavelengths, with no phase offset.

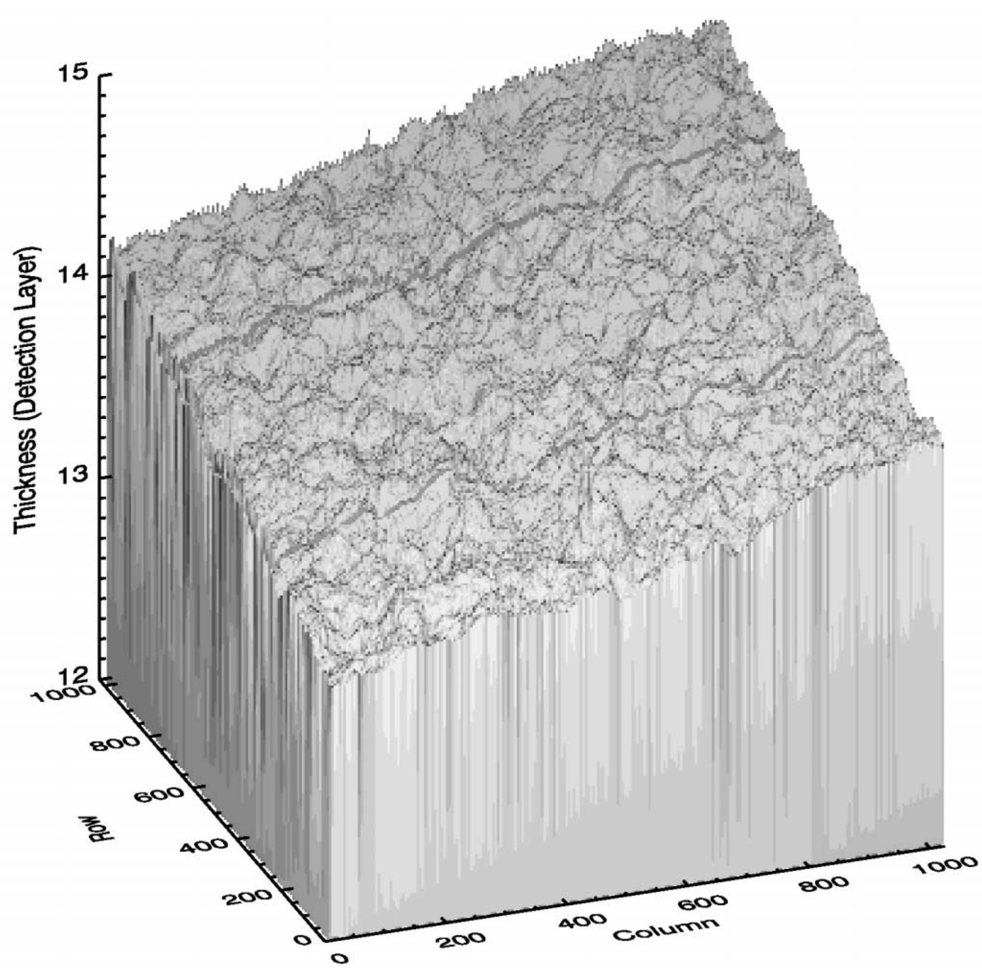

FIG. 12.—Shaded surface plot of the derived thickness of the detection layer of the STIS CCD. The detection layer varies in thickness from 13.24 μm in the lower left corner to 14.83 μm in the upper right corner. Although there are moderate local variations, the overall change in thickness is nearly planar. The rows of the CCD that are shadowed by the fiducials have been replicated from the a nearby unshadowed row for appearance only. These rows contain no information; however, they will never have real spectral data either.

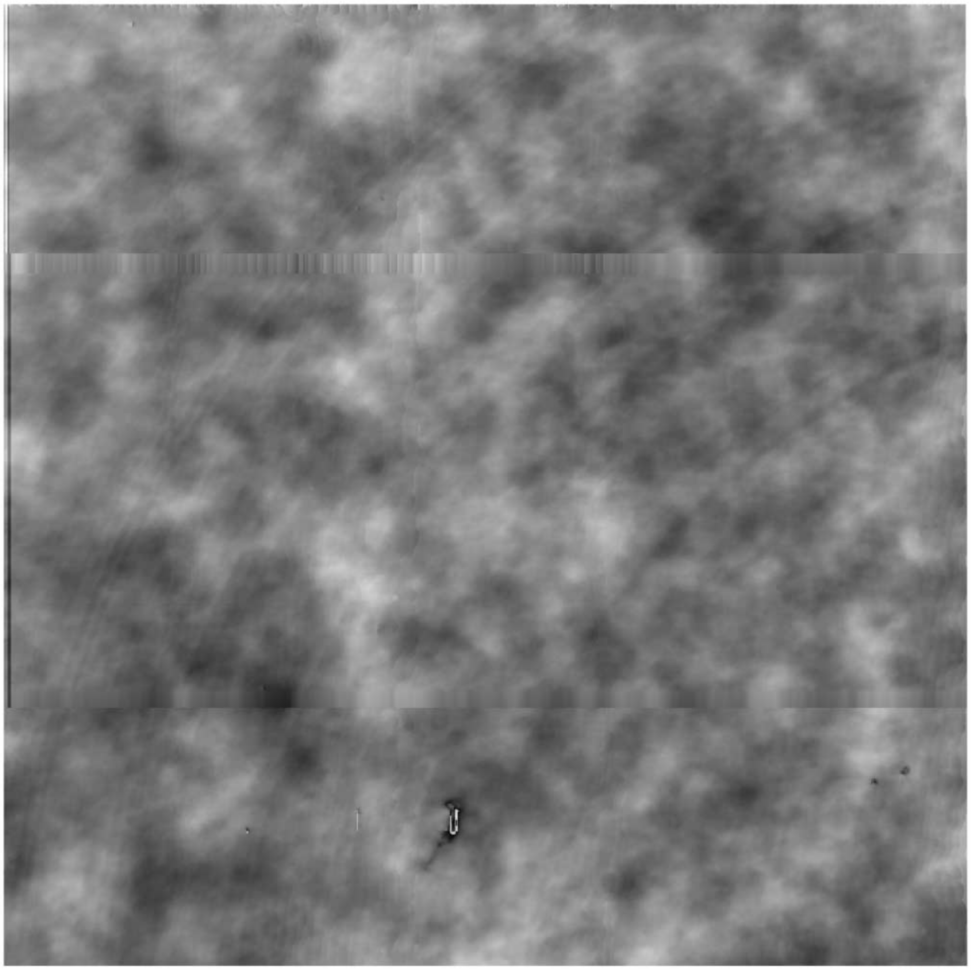

FIG. 13.—Image showing the thickness variations of the STIS CCD detection layer after a plane has been subtracted. The range in thickness is ± 0.07 μm. A regular arc pattern is centered near the lower right corner of the image. In this figure, as in Fig. 12, the rows shadowed by the fiducials have been replicated from a nearby unshadowed row.

in the short-wavelength regime to derive d_3. For those pixels, the gray line is identical to the black line, and there is no information content. Notice how the two lines intertwine, with the gray line being below, then above, the black line in a regular pattern. This is more easily seen in the lower panel of Figure 8, where the difference between the lines is shown. The regularity of the pattern is indicative of some fundamental and systematic difference.

The difference is clarified when we note that the wavelengths of the fringes being fitted vary periodically as we move along the row. The differences in thickness derived from different wavelength regions for the same pixel are clearly a function of wavelength. Of course, the thicknesses of the CCD layers do not change as a function of wavelength, but the refractive index of silicon does.

In reality, we do not solve for the thickness d_3 but rather for the product $n_3 d_3$, where we believe that we know n_3 as a function of wavelength. If instead we assumed that we knew what d_3 was (e.g., from fitting in the long-wavelength regime), then we could solve for n_3.

In Figure 9, we plot the percentage difference in the index of refraction of silicon as a function of wavelength for the wavelength in the low-wavelength regime ($\lambda \leq 8200$ Å) as the open boxes and in the middle-wavelength regime (8200 Å $\leq \lambda \geq 8900$ Å) as the dots. Using these differences, we can adjust the nominal values of the index of refraction of silicon to obtain a revised index of refraction to be used in our model. Note that this does not imply that these are the true values of the index of refraction of silicon or even of the STIS CCD chip's detection layer. To derive these values, we arbitrarily assumed the index of refraction beyond 9000 Å to be accurate, leading to a proper thickness of the detection layer. Our derived *variation* in index of refraction shortward of 9000 Å should be quite good, although the absolute values may not be. Figure 10 shows the revised index of refraction of silicon (*gray line*) and the published index of refraction (HOC III; *dashed line*). Note that the

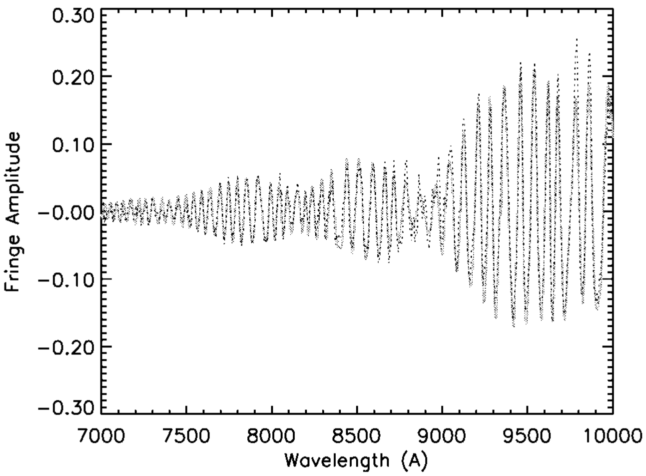

FIG. 14.—Model of rows 511–513 of the STIS CCD is shown as the thick gray line. The observed fringes for those rows (cf. Fig. 1) are shown as the dotted line. Although the amplitudes are not perfect, the shape and spacing of the fringes are well matched.

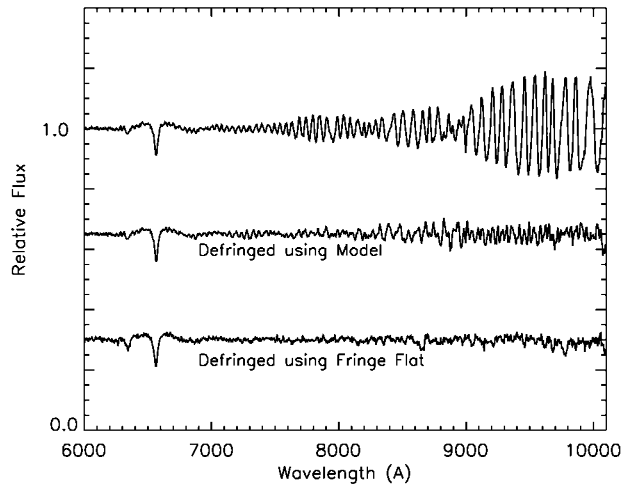

FIG. 15.—Calibration spectrum of the white dwarf standard G191-B2B with the low-frequency spectral shape removed is shown at the top. Each tick mark represents a change of 10% from the mean. A version defringed using our model is plotted offset by −0.35 from the top spectrum. A version defringed using an observed fringe flat is shown offset by −0.70. The observed flat corrects the fringes slightly better than the model. At the high-wavelength end, the model reduces the peak-to-peak fringes from 20% to about 4%, while the fringe flat reduces the fringing to about 3.3%

maximum difference is about 0.4%. This small revision in the value of the index of refraction as a function of wavelength has a profound effect on our ability to model the fringe pattern for real data. Also shown in Figure 10 is the absorption coefficient (k) of silicon from HOC III. Notice how the absorption coefficient becomes quite small at longer wavelengths, giving rise to the behavior seen in Figure 3.

Figure 11 shows the data for the same pixels ($X = 900$–902, $Y = 499$) as presented in Figure 7, but now compared to a model using the revised index of refraction. The model is now in phase at all three wavelengths. Although this revision was derived from fits done in row 499, all pixels all over the CCD behave in the same way (a single value of d_3 now works for all wavelengths) regardless of the wavelength distribution on the pixel.

Given this revised table of the index of refraction of silicon, we proceeded to fit all of the pixels in all of the rows of the CCD (excluding a 5 pixel border and the small regions shadowed by fiducials on the STIS slit) for a single value of d_3 per pixel using the data at all wavelengths between 6500 and 10200 Å. The result of this fitting is a map of the thickness of the detection layer of whichFigure 12 shows a shaded surface plot. The overriding feature is that the detection layer is wedge shaped, varying in thickness from 13.24 μm in the lower left corner to 14.83 μm in the upper right corner.

Figure 13 is an image of the thickness variations after a tilted plane is subtracted from the thickness map. In addition to the general variations in the thickness, there is a series of concentric arcs centered near the lower right corner. Similar marks can be seen in the medium-resolution (G750M) fringe flats. In fact, such arcs are present in many astronomical CCDs. They are due to the manufacture of the CCD chip and may be residual growth rings produced in the silicon boole while it is grown or scoring marks produced while polishing the CCD. That we see them clearly in the thickness map shows that they are, in fact, slight differences in the thickness of the silicon detection layer.

Putting all of these pieces together, we can now make a fringe model for any wavelength distribution across the STIS CCD.Figure 14 shows a model of the central three rows of the CCD (i.e., those shown in Fig. 1) as the thick gray line. The dotted line is from one of the fringe flats taken at the nominal MSM position (cf. Fig. 1). The amplitudes, locations, and wavelengths of the fringes and the nodes are all the consequences of the model and of the physical description of the CCD. The match to the data is quite good: subtraction of the model from the data leaves a maximum residual of ∼4% peak to peak, down from a maximum amplitude of ∼20% peak to peak. For comparison, standard flat fields taken with a narrow slit result in a maximum residual of about 3.3% peak to peak.

4. USING THE MODEL TO DEFRINGE REAL STIS DATA

The true test of the utility of this model is to apply it to real STIS CCD G750L data. In this section, we first use the model to defringe calibration data, in order to evaluate how well the model works relative to the normal calibration procedure. We also use a calibration data set in which the same star was observed at nine different locations within the field to evaluate how the model works as a function of position within the field. We then apply the model to slitless spectra obtained as part of

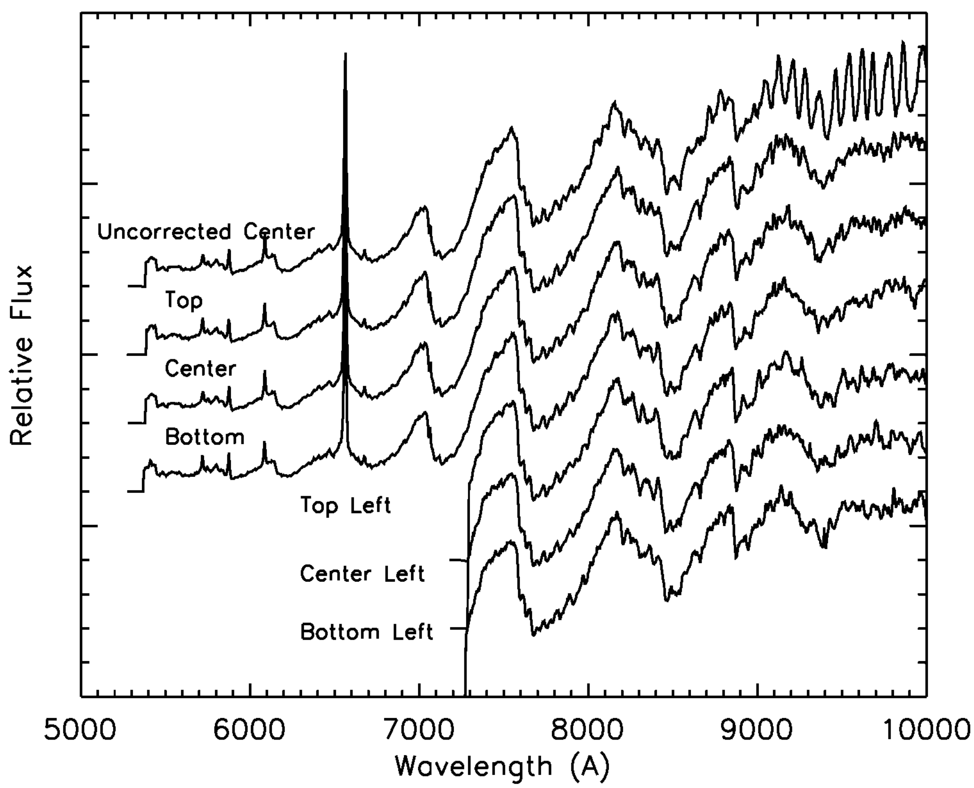

FIG. 16.—Uncorrected STIS spectrum of AX Per obtained near the center of the CCD is shown on the top, while the model defringed spectra obtained at six different locations are shown below. In all cases, the model does a good job in defringing the observed spectrum. There is a slightly higher residual for spectra near the left edge of the detector than for those near the nominal slit location.

the STIS parallel program and evaluate how well the model allows us to pull information from these spectra.

4.1. Calibration Data

4.1.1. G191-B2B

The white dwarf G191-B2B has a relatively featureless spectrum with the exception of the hydrogen Balmer lines. This fact makes it useful as a flux calibration standard for STIS and other spectrographs. Thus, there are a number of observations of G191-B2B taken with the G750L grating. These observations were taken with the star in the slit in the nominal position; the observed spectra can thus have their fringes removed using a fringe flat taken at the time of the observation.

In Figure 15, we show one such spectrum of G191-B2B with the overall shape of the white dwarf continuum removed. On the top, we show the spectrum as reduced except for the presence of the fringes. In the middle is this observation defringed using the model derived here, and at the bottom is the same spectrum defringed using a fringe-flat image. The model does well in defringing this featureless star located on the slit. The residuals have been reduced by about a factor of 5, from 20% peak to peak (10% rms) to 4% peak to peak (2% rms). However, the model does not do quite as well as the actual calibration fringe flat, which reduces the fringing to 3.3% peak to peak (1.7% rms). For comparison, the rms noise due to photon statistics is only about 0.5% (~40,000 photons pixel^{-1} in the 8000–10000 Å wavelength range).

4.1.2. AX Per

The experiment with G191-B2B shows that we can get good fringe models for stars that are in the nominal position (i.e., on the slit). However, that might be a favored position, since the data used to derive the model were all taken through the slit. We cannot make a direct comparison with fringe removal using a fringe flat at any other location on the CCD. There is, however, a data set in which a spectrum of the bright star AX Persei was obtained at nine different positions in the field. These data were obtained to calibrate the wavelength scale for slitless spectra, with the star in a 3×3 raster of positions through the top, middle, and bottom of the CCD (rows 120, 515, 910 and columns 120, 515, 910). The three spectra along the right-hand side of the CCD are not of use for checking defringing because the fringing wavelengths fall outside of the detector. The other six locations are of use, however. Figure 16 shows the uncorrected spectrum of AX Per obtained in the center (pixel $X = 515$, $Y = 515$) location, as well as the six

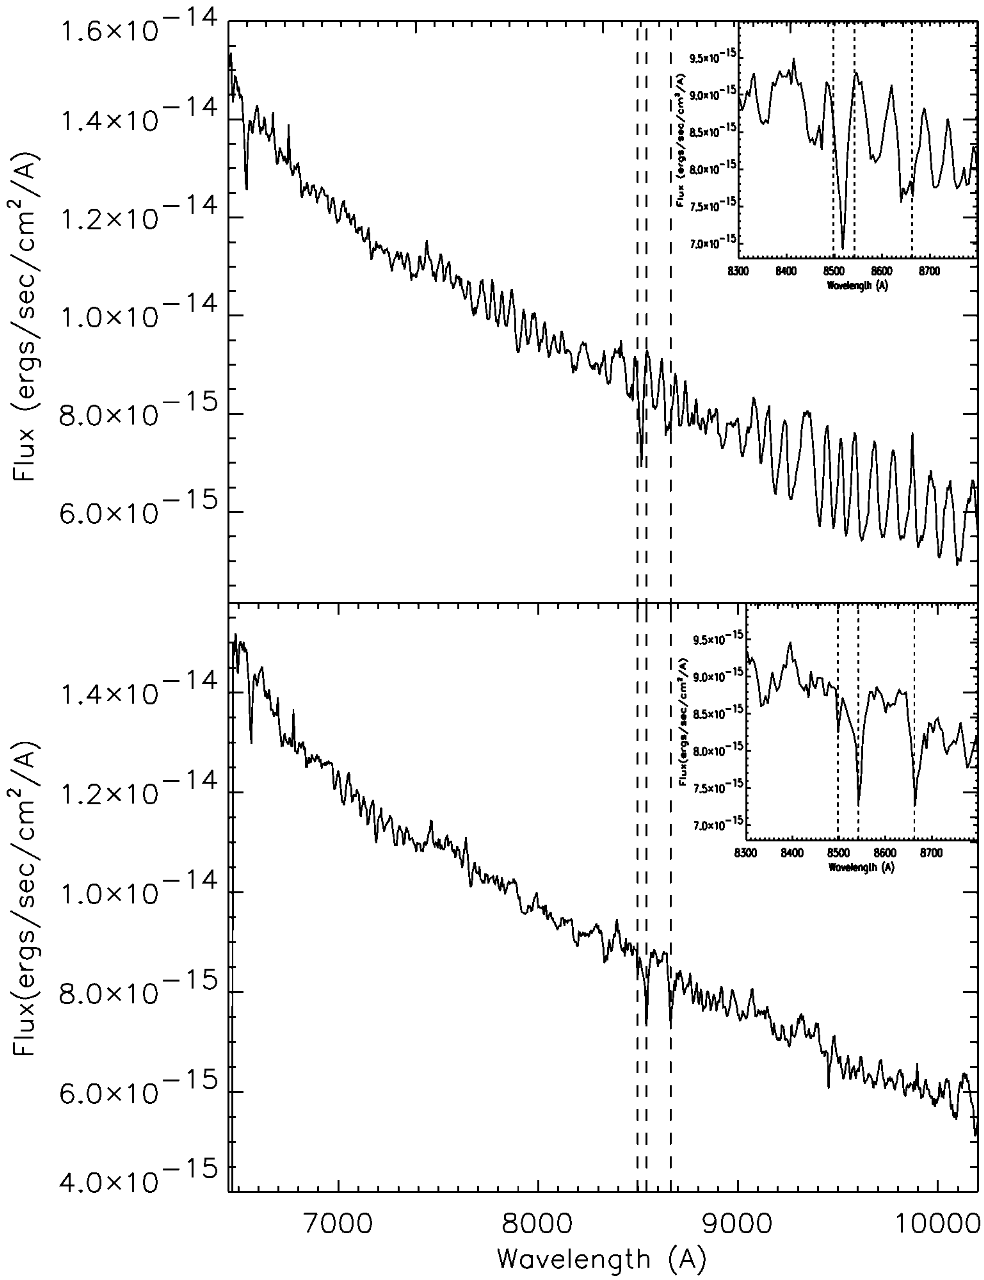

FIG. 17.—Defringing of a slitless spectrum is demonstrated in these raw (*upper*) and defringed (*lower*) plots of a bright F star found in one of the STIS parallel program fields. In addition to the general improvement at wavelengths longward of 9000 Å, the Ca II triplet lines (rest wavelengths indicated by the vertical dashed lines) are much more distinct and apparent (see inset). Also, note the improved wavelength scale, as determined by the λ8542 line. The photon noise for this spectrum is ~0.5% at $\lambda = 9500$ Å.

left-hand side (column 120) and center (column 515) locations corrected for fringing using the model. At all locations, the defringed spectrum shows a marked improvement over the uncorrected spectrum. Close inspection shows that in the high fringe amplitude region (~9000–10000 Å), the three spectra taken at the left-hand side of the CCD have slightly higher residuals. The rms fringe amplitude in this wavelength region is on average 10.9% before defringing. The average rms residual after defringing the data obtained on the central column is 2.6%, while the average rms residual of the three spectra obtained on the left (column 120) is 3.2%. Not surprisingly, the model works slightly better near the slit location; however,

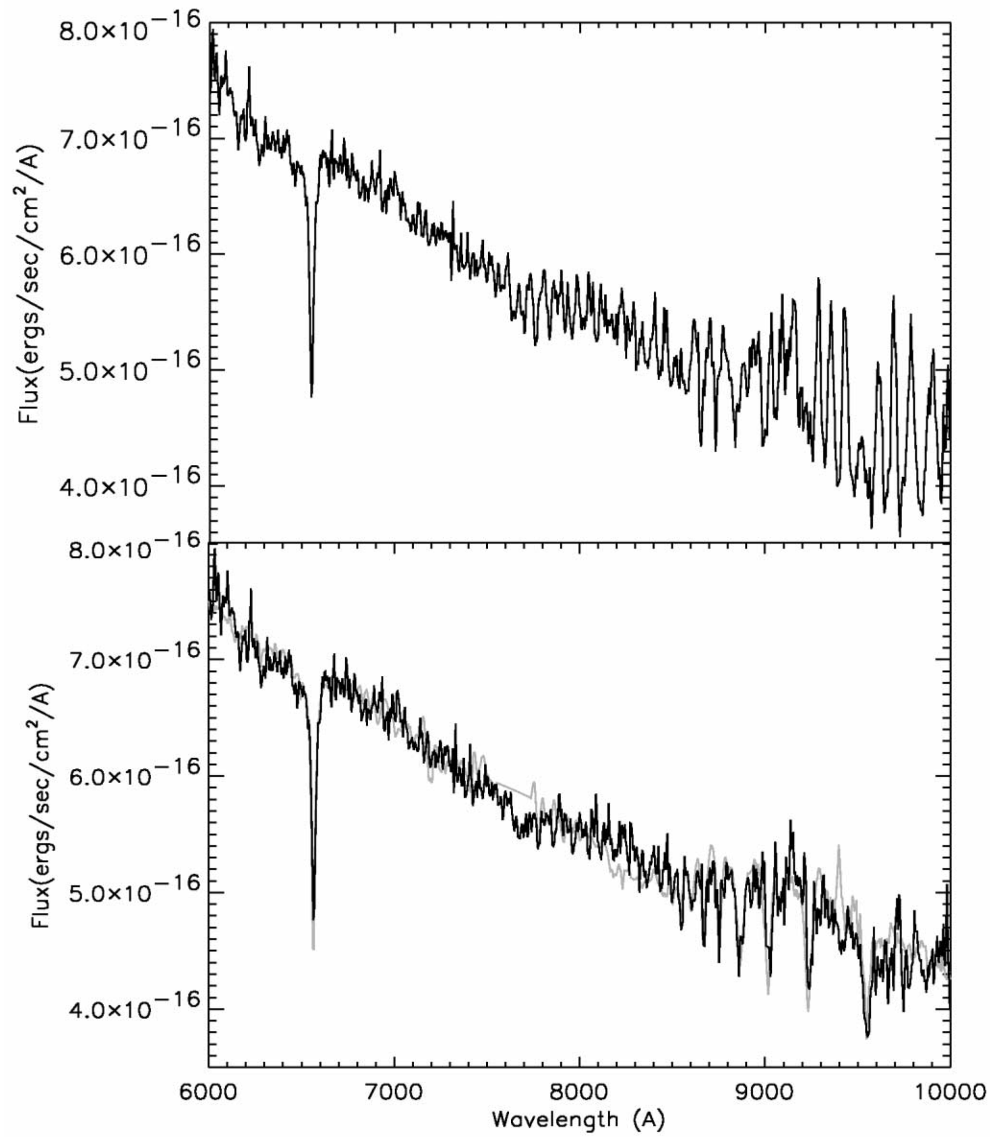

FIG. 18.—Raw and corrected plots of a bright, reddened A star from the STIS parallel observations. A reddened spectrum of an A5 V star (Pickles 1998) is shown as the thick gray line in the lower (defringed) plot. Note the Paschen series longward of 8500 Å. The photon noise for this spectra is approximately ∼2.5% at $\lambda = 9500$ Å.

the fringe amplitudes are reduced by a factor of 3–4 elsewhere in the field. The rms photon noise for the AX Per data is ∼0.8%.

4.2. Parallel Data: Slitless Spectra

Since 1997 June 2, the STIS has been taking parallel images and slitless spectra of random locations on the sky whenever another instrument is prime (Gardner et. al. 1998). Until now, there has been no way to calibrate these spectral images for the effects of fringing. The true test of our model is to see how well we can reduce the fringing in high signal-to-noise ratio stars in the parallel fields.

4.2.1. A Strategy for Defringing Extracted Spectra

At first glance, once the model is calibrated, defringing a spectrum should be easy. Find the object on the STIS field image, use that location to define the wavelength scale, extract the spectrum; run a model for those pixels, normalize the fringe model, and divide the extracted spectrum by the normalized model. Unfortunately, we do not know the absolute wavelength scale a priori to the accuracy required, as a result of thermal drift and MSM nonrepeatability. A few-pixel shift in the wavelength scale can cause a phase shift in the fringe pattern and hence poor defringing.

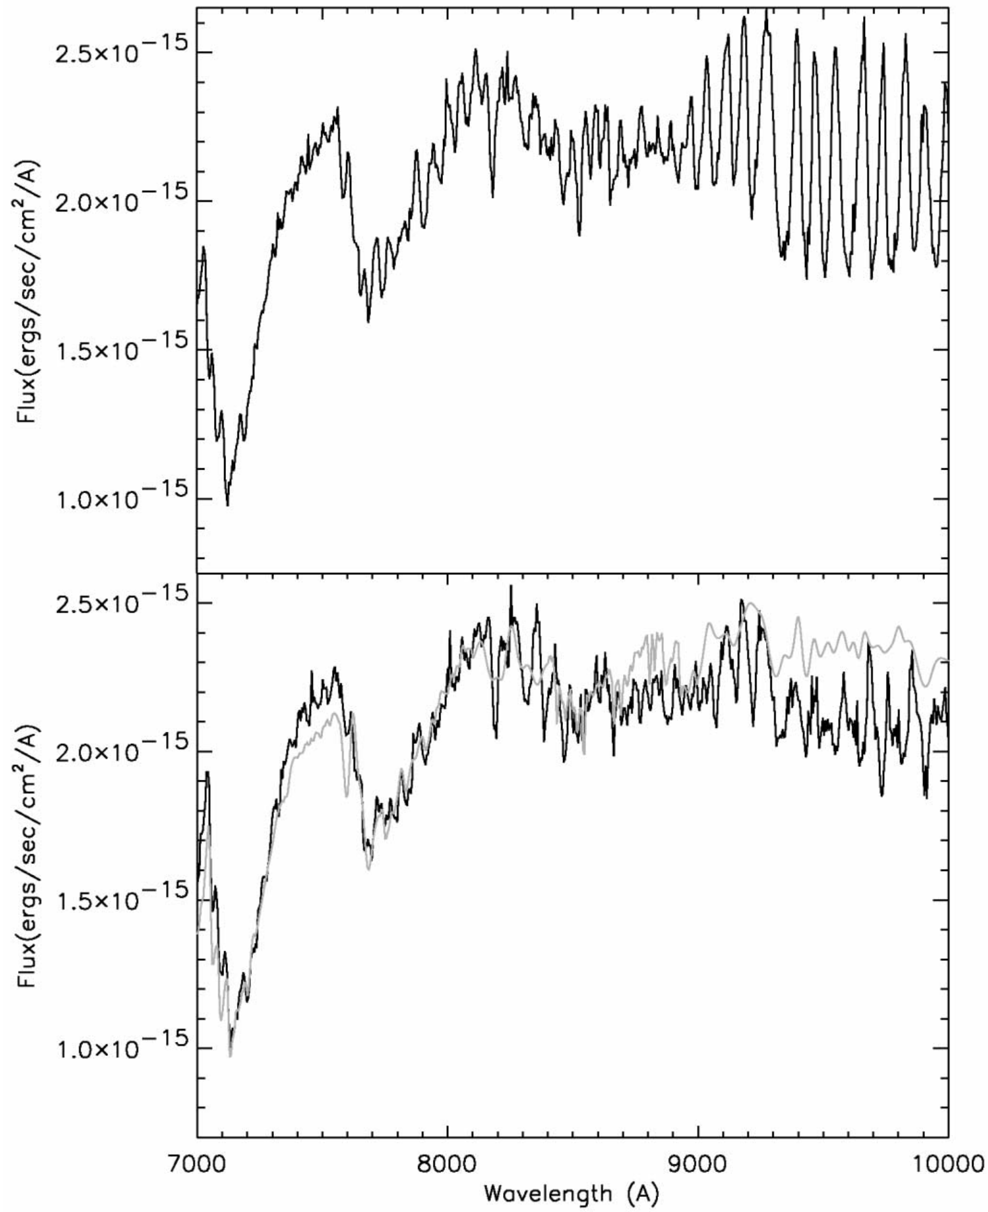

FIG. 19.—Raw and corrected plots of a bright M star from the STIS parallel observations. The thick gray line in the lower (defringed) plot is a Pickles (1998) spectrum of an M3 V star. The photon noise for this spectra is ∼1.2% at $\lambda = 9500$ Å.

We have, therefore, devised a strategy of defringing the spectrum using a number of slightly different trial wavelength scales and measuring the residual modulation in various regions of the resulting defringed spectra. The wavelength scale that minimizes the residual modulation is determined to be the correct wavelength scale, leading to the most accurate defringed spectrum. As a by-product, this procedure yields a sensitive determination of the wavelength zero point, which is especially useful for objects not located near the nominal slit position. The resulting accuracy is essentially that of the wavelength scale for the 51 flat-field calibration images used to create the model. The error in this determination, which is made using observations of a slit-filling emission-line lamp, is negligible compared to that in the wavelength zero point for an actual star. Indeed, the sensitivity of the fringe location to wavelength is such that fringing can be used to improve the wavelength zero point even for slit spectroscopy, in which zero-point shifts can result from poor centering of the star in the slit.

4.2.2. Results

To illustrate the usefulness of this modeling, we will consider three stars chosen from the many stars in the parallel archives

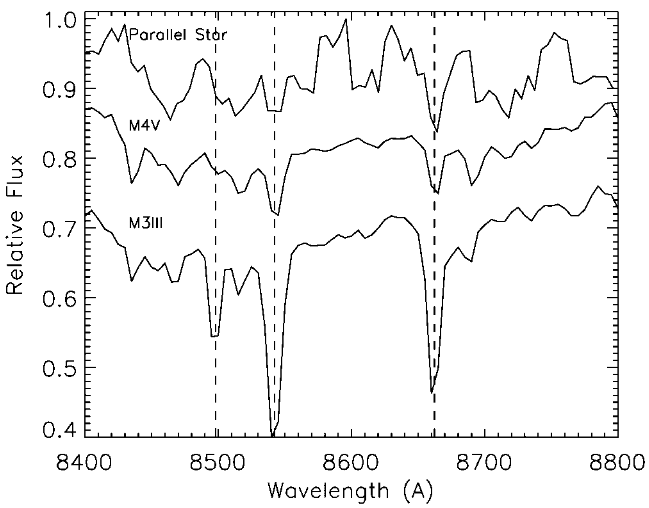

FIG. 20.—Region of the Ca II triplet is shown for the STIS parallel spectrum in Fig. 19. Comparison spectra of an M4 main-sequence star and an M3 giant star (Silva & Cornell 1992) are also shown, offset by −0.12 and −0.24, respectively. The M giant has much stronger Ca II lines than the main-sequence star. Defringing allows us to see that the STIS parallel star has weak Ca II lines, consistent with its being a main-sequence star.

that have both well-exposed images and slitless spectra. The three stars are located on different regions of the CCD and have very different spectra. One is most likely an early- to mid-F star. Another is a reddened mid-A star, and the third is an M star. In each case, the fringing affects our ability to match the spectrum to model spectra in a different way.

4.2.2.1. An F Star

Figure 17 presents the spectrum of a star at image position $X = 298$, $Y = 358$, relatively far from the nominal center of the CCD. The star is in the constellation Monoceros at an R.A. = $7^h24^m23^s$, decl. = $-00°27'13''$ (J2000.0).

Figure 17 shows the spectrum of this star before and after defringing. The inset at the upper right of each panel shows the area around the near-IR Ca II triplet. While the amplitude of the fringes in this wavelength region is only about 8%, it is still enough to mask two of the three Ca II lines and alter the measured strength of the third. In the upper panel (not defringed), there is a dip near where the strongest of the Ca II lines (λ8542) coincides with a fringe minimum. As described above, the wavelength scale is uncertain for objects away from the center of the field. From this spectrum, one might suspect that the dip was the λ8542 line, but one would not be able to say how strong the line is, since much of its depth is due to the fringing. The other Ca II lines are indistinguishable from fringes. In the lower panel, the wavelength zero point has been shifted in the process of producing the defringed spectrum. All three Ca II lines are now apparent, and the fringing has been reduced throughout the spectrum. This result can be compared to the spectral library of Pickles (1998). We find that this star is most likely an F2–F5 main-sequence star with a *V* magnitude of ∼13.2 at a distance of about 900 pc.

4.2.2.2. An A Star

Figure 18 shows the raw and defringed spectra of a star at pixel location $X = 525$, $Y = 841$ in another parallel image. In this spectrum, the fringing seems to be less regular in the 9000–10000 Å range. This is due to the combination of severe fringing at the ∼20% level and fairly strong absorption lines. In the lower panel, defringing reveals the Paschen series. A reddened Pickles spectrum of an A5 V star is shown as the thick gray line. This star, at R.A. = $5^h41^m43^s$, decl. = $39°10'57''$ (J2000.0), is in the constellation of Auriga at a Galactic longitude of 170° and latitude of 5°. An approximate *V* mag of 16.6 with the assumed reddening implies a distance of ∼4 kpc.

4.2.2.3. An M Star

Figure 19 shows the raw and defringed spectra of a star located in a yet another part of the CCD (pixel $X = 184$, $Y = 978$). It is a very different spectrum from the other two, in that it is an M3 main-sequence star. The lower panel shows the defringed spectrum overlaid on Pickles's spectrum of an M3 V star. The match is quite good, except that the red end of the continuum is lower than the model. Note that the Ca II lines are quite weak; this is an important feature in distinguishing between a main-sequence M star and an M giant, as shown in Figure 20, a plot of the region around the Ca II triplet. The top spectrum is the spectrum of the star in the parallel field, the middle spectrum is an M4 main-sequence star (Silva & Cornell 1992), and the lower spectrum is an M3 giant (Silva & Cornell 1992). The strengths of the Ca II lines are much more like those in the M4 main-sequence star than those in the M3 giant. This would be impossible to tell in the presence of the fringes. This star, at R.A. = $4^h20^m50^s$, decl. = $56°12'05''.7$ (J2000.0), is in the constellation of Camelopardus at a Galactic longitude of 149° and latitude of 4°; a *V* mag of ∼16.86 implies a distance of ∼200 pc.

5. SUMMARY

We have presented a method that allows us to model the effects of fringing in the STIS CCD. This method should be transportable to other CCDs.

We have shown that with sufficient calibration data, the physical structure of the CCD can be solved for as a function of position. Armed with this information, we can produce a model of the fringing for any wavelength distribution on any pixel.

We have solved for the physical structure of the STIS CCD for most of the 1024 × 1024 pixels (excluding the edges), and in doing so, we revised the index of refraction of silicon as a function of wavelength in order to achieve a consistent solution.

We have devised a strategy to use this physical solution and model function to make a "fringe flat" for slitless spectra of objects anywhere in the STIS field of view.

We have evaluated the effectiveness of the model in defringing STIS calibration spectra. The peak fringe amplitude of G191-B2B is reduced from 20% to 4% peak to peak, whereas defringing through the normal means reduces the fringe amplitude to 3.3% peak to peak. Observations of AX Per obtained in six different locations show that the model will reduce the fringe amplitudes at a variety of positions, although it does slightly better near the nominal slit location.

We have demonstrated the utility of this model by defringing the spectra of three very different stars observed on three very different areas of the CCD. These stars were all in STIS parallel observations, and all had important stellar features that were fully or partly obscured by the fringing. The model defringing reduced the effects of the fringes by a factor of 4–5 (from ∼20% peak to peak to ∼4% peak to peak) allowing us the see and match these important lines to observed spectra of the same stellar type.

Instrument builders can enhance the accuracy of CCD observations by taking fringe spectra (or a series of monochromatic images) in the laboratory. The data set used in this paper was the minimum necessary because of the use of on-orbit time. Even though ground-based laboratory time is also at a premium even for ordinary purposes such as flats, biases, and wavelength calibrations, we suggest that a careful exploration of the fringing is desirable for CCDs to be used at long wavelengths. For the STIS, the f/36 input beam allowed us to use the normal incidence approximation. We suggest that even for CCDs intended to be used in fast beams, laboratory data should be obtained in a slow beam to permit this simplification of the analysis. Once the physical properties of the thin film stack have been determined, the modeling of fringing in a fast beam should be done using a more sophisticated algorithm that properly represents the distribution of incident angles and wavelengths, as well as the more complicated interference inside the chip. Whenever the problem can be approximated as a monochromatic plane wave encountering a thin film stack of infinite extent, the IMD program (Windt 1998) or a similar algorithm can be applied.

REFERENCES

Born, M., & Wolf, E. 1999, Principles of Optics (7th ed.; Cambridge: Cambridge Univ. Press)
Gardner, J. P., et al. 1998, ApJ, 492, L99
Janesick, J. R. 2001, Scientific Charge-Coupled Devices (Bellingham: SPIE Press)
Kimble, R. A., et al. 1994, Proc. SPIE, 2282, 169
Kimble, R. A., et al. 1998, ApJ, 492, L83
Palik, E. D. 1998, Handbook of Optical Constants of Solids III (Sydney: Academic Press)
Pickles, A. J. 1998, PASP, 110, 863
Silva, D. R., & Cornell, M. E. 1992, ApJS, 81, 865
Windt, D. L. 1998, Comput. Phys., 12, 360