

Spextool: A Spectral Extraction Package for SpeX, a 0.8–5.5 Micron Cross-Dispersed Spectrograph

MICHAEL C. CUSHING,^{1,2}

Institute for Astronomy, University of Hawai‘i, 2680 Woodlawn Drive, Honolulu, HI 96822; mcushing@arc.nasa.gov

WILLIAM D. VACCA¹

NASA Ames Research Center, MS 144-2, Moffett Field, CA 94035; and Department of Astronomy, University of California, 601 Campbell Hall, Berkeley, CA, 94720; wvacca@mail.arc.nasa.gov

AND

JOHN T. RAYNER

Institute for Astronomy, University of Hawai‘i, 2680 Woodlawn Drive, Honolulu, HI 96822; rayner@irtf.hawaii.edu

Received 2003 May 9; accepted 2004 January 15; published 2004 April 2

ABSTRACT. We describe an IDL-based package for the reduction of spectral data obtained with SpeX, a medium-resolution, 0.8–5.5 μm cross-dispersed spectrograph and imager for the NASA Infrared Telescope Facility. The package, called Spextool, carries out all the procedures necessary to produce fully reduced spectra including preparation of calibration frames, processing and extraction of spectra from science frames, wavelength calibration of spectra, and flux calibration of spectra. The package incorporates an “optimal extraction” algorithm for point-source data and also generates realistic error arrays associated with the extracted spectra. Because it is fairly quick and easy to use, requiring minimal user interaction, Spextool can be run by observers at the telescope to estimate the signal-to-noise ratio of their data. We describe the procedures incorporated into Spextool and show examples of extracted spectra.

1. INTRODUCTION

Cross-dispersed spectrographs provide one of the most efficient means of acquiring spectroscopic data from celestial sources, because they record a broad wavelength range in a single exposure. Concomitant with the increase in efficiency, however, is a substantial increase in the complexity of the methods needed to reduce these data. Usually the spectral orders are curved and irregularly spaced on the detector so that the dispersion axis is not aligned with the detector rows. In addition, the spatial axis is sometimes not aligned with the detector columns, so that a given wavelength may span many columns along the slit. These properties can make the extraction and calibration of cross-dispersed spectra particularly challenging.

Until recently, cross-dispersed spectrographs have been mainly limited to optical wavelengths. However, with the advances in near-infrared detector technology, a new generation of cross-dispersed spectrographs sensitive in the 1–5 μm range has

been developed. One such instrument is SpeX (Rayner et al. 2003), a medium-resolution ($R \equiv \lambda/\Delta\lambda \sim 2000$), 0.8–5.5 μm spectrograph operating at the 3.0 m NASA Infrared Telescope Facility on Mauna Kea.

In an effort to provide a set of semiautomatic routines for reducing SpeX data, we have developed a spectral extraction and reduction package called Spextool (spectral extraction tool). Written in the interactive data language (IDL), this package is driven by a set of graphical user interfaces (GUIs) in which a user specifies various parameters pertinent to the reduction procedures and performs the reductions with a few mouse clicks.

In this paper, we present a description of Spextool and provide examples of spectra reduced with Spextool. In § 2 we briefly describe the aspects of SpeX that relate to the data reduction process. In § 3 we describe the typical observing modes used with SpeX. In § 4 we provide a detailed description of the reduction procedures carried out by Spextool, and in § 5 we give a summary.

2. SpeX

A full description of the design, construction, and performance of SpeX is provided by Rayner et al. (2003). Here we

¹ Visiting Astronomer at the Infrared Telescope Facility, which is operated by the University of Hawai‘i under contract from the National Aeronautics and Space Administration.

² Current address: SETI Institute, NASA Ames Research Center, MS 245–3, Moffett Field, CA 94035.

TABLE 1
SPEX OBSERVATION MODES

Mode	Wavelength Range (μm)	R^a	Slit Lengths (arcsec)
Short XD	0.8–2.4	200–2000	15
Long XD1.9	1.9–4.2	250–2500	15
Long XD2.1	2.1–5.0	250–2500	15
Long XD2.3	2.3–5.4	250–2500	15
Low Resolution	0.8–2.5	25–250	15, 60
Short Single Order	0.8–2.5 ^b	200–2000	15, 60
Long Single Order	1.9–5.5 ^b	250–2500	15, 60

^a R is the resolving power and is given by $R = \lambda/\Delta\lambda$; the two values correspond to the widest ($3''$) and narrowest ($0''.3$) slits, respectively.

^b Requires multiple grating settings to achieve full wavelength coverage.

will discuss only those aspects of the instrument that are relevant to the data reduction process.

Table 1 provides a summary of the various observing modes for SpeX. Spextool was designed primarily to extract spectra from the cross-dispersed modes, since single-order spectra can be extracted easily using other data reduction packages (e.g., IRAF). The short-wavelength cross-dispersed (SXD) mode of SpeX covers the wavelength range 0.8–2.4 μm with an average resolving power $\langle R \rangle = 2000$ for a slit width of $0''.3$. The entire wavelength range, except for a 0.6 μm gap in the telluric band between the H and K bands, is recorded simultaneously in 6 orders. As can be seen in Figure 1, the spectral orders are substantially curved, and the spacing between the orders varies with position on the array. The long-wavelength cross-dispersed (LXD) mode covers the wavelength range 1.9–5.5 μm with $\langle R \rangle = 2500$ for a slit width of $0''.3$. The entire wavelength range, except for a 0.1 μm gap centered at 4.4 μm , is covered in 7 orders, but only 6 orders can be recorded simultaneously on the detector. Therefore, the LXD mode actually consists of three submodes (LXD1.9, LXD2.1, and LXD2.3) that cover slightly different wavelength ranges (see Table 1). Both cross-dispersed modes have a slit length of $15''$, which corresponds to ~ 100 pixels on the detector. Spextool also has the ability to reduce data taken in the LOWRES (prism) mode with the $15''$ long slit (~ 100 pixels); this mode covers the wavelength range ~ 0.8 –2.5 μm at $\langle R \rangle = 250$ for a slit width of $0''.3$. Wider slits ($0''.5$, $0''.8$, $1''.6$, and $3''.0$) are also available for each mode and provide proportionally lower resolving powers. The pixel size for all three modes is $0''.15$ pixel⁻¹, so even during the best seeing conditions currently attainable at the IRTF ($\sim 0''.4$ at 2.2 μm), the point-spread function is sampled by ~ 3 pixels. As a result, intrapixel sensitivity variations (e.g., Lauer 1999) are not a concern for the data reduction process.

The gratings are operated in the near-Littrow configuration, which means, in principle, the image of the slit should not be tilted with respect to the detector columns (Schroeder 2000). In practice, however, the image of the slit is tilted by up to ± 0.5 pixels over the length of the $15''$ slit because of pre-dispersion from the cross-dispersing prisms, the natural curvature




FIG. 1.—Raw SXD flat-field exposure. The spatial axis is parallel to the columns and wavelength increases to the right. The six cross-dispersed orders (3–8; bottom to top) are apparent. The diagonal lines are 1 pixel wide cracks in the array. The spot at the approximate coordinates (600, 275) is a cluster of bad pixels. The faint interorder flux is due to scattered light and internal reflections (“ghosts”) and is typically at a level of $\sim 1\%$ of the signal in the brightest order.

due to any finite-length slit, and the precision to which the slits can be aligned with the detector columns (see Rayner et al. 2003). Since this misalignment is significantly smaller than the maximum spectral resolution delivered by SpeX, we have chosen not to attempt to correct for it in Spextool. This greatly simplifies the data reduction process, since the spectral data do not have to be rectified before extraction.

An external calibration unit contains lamps needed for flat-fielding and wavelength calibration. The flat-field lamps are a 10 W quartz-tungsten-halogen lamp (3200 K) and 0.1 W bulb for the 0.8–2.5 μm wavelength range, and an infrared source (1100 K) for the 2.5–5.5 μm wavelength range. The unit also contains a low-pressure argon discharge lamp for wavelength calibration. The lamps are located at the entrance ports of an integrating sphere. Light from the exit port of the integrating sphere is projected into the spectrograph at the same focal ratio as light from the telescope.

3. OBSERVING TECHNIQUE

It is instructive to discuss the typical manner in which observations are acquired with SpeX, as the observing procedures dictate, in part, the data reduction process. The number of counts recorded in any given detector pixel is given by the sum of the counts from the dark current, emission from the telescope and sky, and the target. This is an ideal case that ignores any systematic effects such as array persistence (see Vacca, Cushing, & Rayner 2004). In order to remove the dark current, telescope signal, and sky signal (to first order) from the data, observations are usually taken in pairs: an **A** frame exposure on the target immediately followed by a **B** frame exposure of

equal integration time to measure the dark current and the telescope and sky background alone (e.g., Joyce 1992). (Hereafter we use a bold letter \mathbf{X} to represent an image, and the subscripted designation $X_{i,j}$ to represent the value of the image at the pixel in column i and row j .) Mathematically we have

$$\mathbf{A} = \mathbf{dark} + \mathbf{telescope} + \mathbf{sky}_A + \mathbf{target}, \quad (1)$$

$$\mathbf{B} = \mathbf{dark} + \mathbf{telescope} + \mathbf{sky}_B, \quad (2)$$

where we have explicitly allowed for the possibility that the flux level from the sky may change between the \mathbf{A} and \mathbf{B} exposures (e.g., due to the time-variable OH^- line emission from 0.8–2.5 μm or thermal background from the sky from 2.5–5.5 μm) but have assumed that the non-sky background signals, including the dark current and telescope emission, do not (see, however, Rayner et al. 2003). Subtracting the \mathbf{B} frame from the \mathbf{A} frame (pair subtraction) gives

$$\mathbf{A} - \mathbf{B} = \mathbf{target} + \mathbf{sky}_A - \mathbf{sky}_B \quad (3)$$

$$= \mathbf{target} + \mathbf{R}, \quad (4)$$

where we have defined $\mathbf{R} = \mathbf{sky}_A - \mathbf{sky}_B$ to be the residual sky signal. As long as the signal from the sky in the individual raw \mathbf{A} and \mathbf{B} frames has not reached the saturation level of the detector, any residual sky signal can easily be subtracted in the extraction process (see § 4.2.5).

If the object is an extended source, then the \mathbf{B} frame is taken after offsetting (“nodding”) the telescope to a patch of nearby blank sky so that the object is no longer on the slit (A-sky mode). If the object is a point source, however, a more efficient observing method involves nodding the telescope so that the object is moved to a new (second) position along the slit and then taking the \mathbf{B} frame exposure (AB mode). In this mode, subtracting the \mathbf{B} frame from the \mathbf{A} frame leaves both a positive and negative signal (spectrum) from the object and any residual signal from the sky. In either mode, the strength of the residual sky signal depends on the variability of the sky and the rapidity with which the nodding was performed.

In order to correct for absorption from the Earth’s atmosphere, as well as to assign physical flux units to the detected counts in the extracted object spectrum, observations of a so-called telluric standard star must also be obtained. In addition to the science and telluric standard data, calibration frames that include a series of flat-field exposures as well as argon arc exposures for wavelength calibration are also taken. The object and standard star exposures, along with the flats and arcs, make up a single data set.

4. SPECTRAL REDUCTION

The reduction process consists of three main steps:

1. *Preparation of calibration frames.*—Combine and nor-

malize flat-field exposures and create an arc image suitable for wavelength calibration.

2. *Spectral extraction.*—Extract and wavelength calibrate target and standard star spectra.

3. *Post-extraction processing.*—Correct target spectra for telluric absorption, merge spectra from different orders, clean and smooth spectra, and merge spectra from different modes (SXD and LXD).

A flowchart showing the order of the various steps is shown in Figure 2, and we describe each step in detail in §§ 4.1, 4.2, and 4.3, respectively. The GUI used to perform steps 1 and 2 is shown in Figure 3.

4.1. Step 1: Calibration Frames

4.1.1. Flat Fields

As described in § 3, multiple flat-field exposures are taken as part of the calibration set and then combined to maximize the signal-to-noise ratio (S/N). Before the individual flat-field frames are combined, Spextool computes the median signal level of each frame. It then scales each frame to a common median flux level (the median of the individual median signal levels) to compensate for any variations in the flux levels due to changes in the temperature of the lamp filaments. These temperature changes also affect the slope of the spectrum produced by the filaments. However, we have found that correcting for these changes makes almost no difference (less than 0.1%) in the final normalized flat field, and so in the interest of computational speed, we do not correct for these changes. The scaled flat-field exposures are then combined using the statistic chosen by the user (median, mean, or weighted mean).

One of the most important steps in the entire reduction process is the determination of the location of the orders on the array. The combined flat-field image, with its high S/N across the entire wavelength range for all the orders, is ideal for this purpose. Spextool determines the spectroscopy mode in which the data were taken by reading the FITS header of the data frames and retrieves the approximate center positions of each order from a file stored on disk. Starting from these center positions, an edge detection algorithm moves across the array at regularly spaced intervals and locates the bottom and top edge of each order. The individual edge points are then fitted with a low-order robust polynomial (see Fig. 4, *left*). The resulting *edge coefficients* then provide the location of each order at any column. As can be seen in Figure 4 (*left*), some orders are truncated by the edges of the array. Therefore, in addition to the edge coefficients, the columns at which the truncation occurs are also identified. Spectral extraction proceeds only between the truncation limits, because knowledge of the order location is crucial for the extraction process.

The 15" slit is re-imaged across ~ 100 pixels, which yields an average plate scale of $0''.15 \text{ pixel}^{-1}$. However, because of anamorphic magnification, the slit height measured in pixels

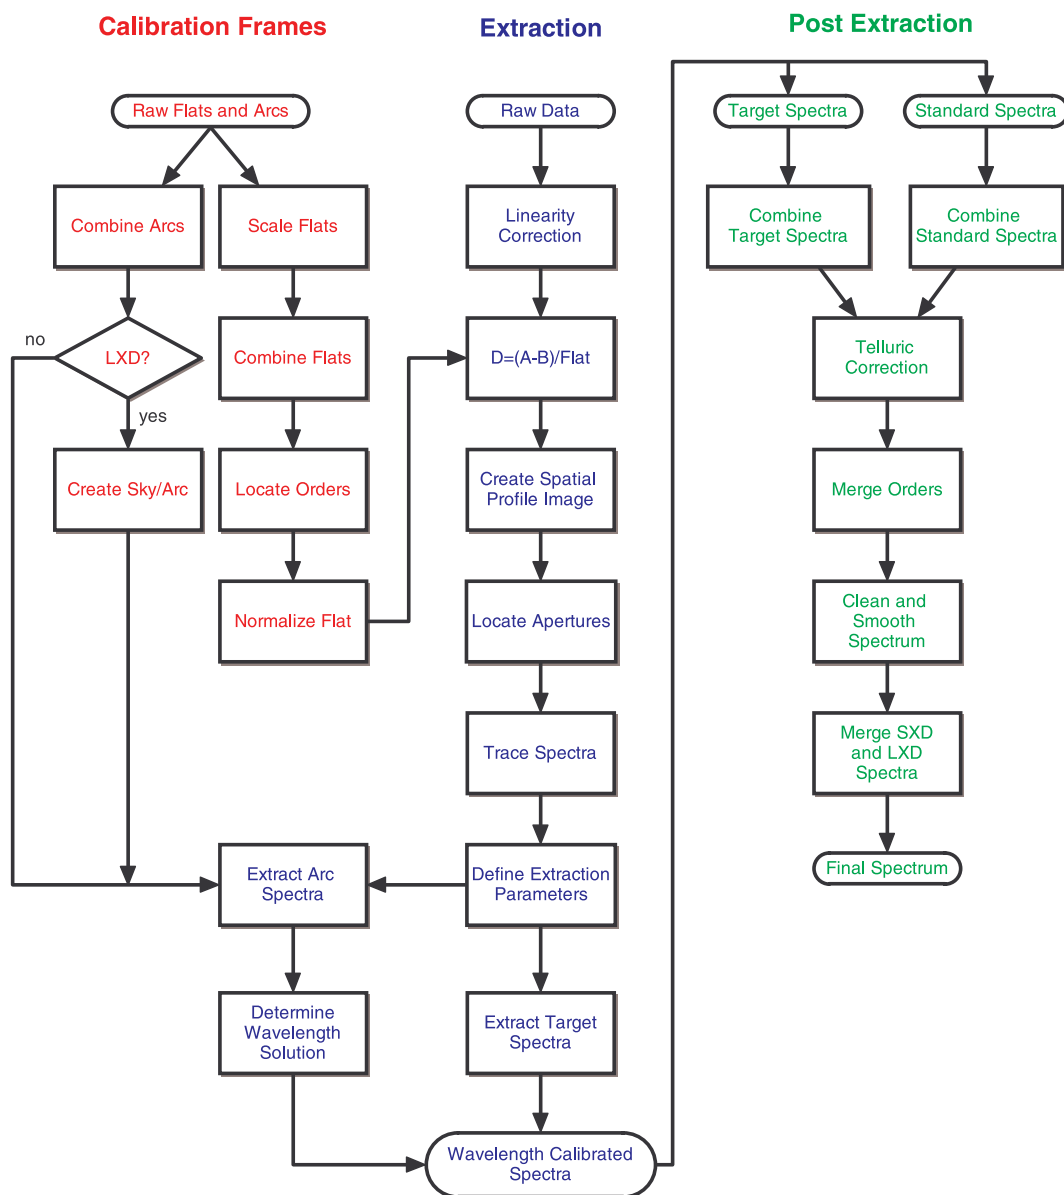


FIG. 2.—Flowchart showing the steps in the reduction process used by Spextool.

varies within an order (≈ 3 pixels) and from order to order. For example, order 3 in the SXD mode has an average slit height of 97 pixels, while order 8 has an average slit height of 106 pixels. To simplify spectral extraction, all operations that involve the spatial dimension are carried out in units of arcseconds instead of pixels. The arcsecond value of each pixel along the slit at a given column is determined by assuming that the length of the slit in pixels is always equal to $15''$.

The flat-field frame must be normalized to unity for users to obtain estimates of the S/N of their spectra. To facilitate the normalization process, each order is “straightened” by resampling it onto a uniform grid of i columns and s spatial points.

Since we assume that the reimaged slit is aligned with the columns of the array, the resampling proceeds on a column-by-column basis. The result is a rectangular image of the order uniformly sampled in the spatial dimension. Because the lamps do not illuminate the slit the same way as the sky does, a two-dimensional surface is then fitted to the straightened order image of the flat. The raw flat field is normalized by resampling the surface fit back onto the arcsecond grid of the raw data column by column and dividing the result into the raw data. It should be noted that the final normalized flat field has not been resampled. The resampling occurs only to construct the smooth, two-dimensional surface fit. The normalized flat field




FIG. 3.—Main Spextool GUI.

is then written to disk with the edge coefficients and the truncations limits contained in the FITS header. Figure 4 (*right*) shows an example of a normalized SXD flat field.

4.1.2. Arcs

Wavelength calibration of SpeX data is accomplished using argon lines. For the SXD and LOWRES modes, several arc exposures are combined to increase the S/N. For the LXD modes, an “arc off” frame (taken immediately after the arc lamp is turned off) is subtracted from an “arc on” frame (taken with the arc lamp turned on) in order to remove the strong thermal continuum from the heated arc lamp.

For $\lambda \gtrsim 3 \mu\text{m}$ the number and strength of the argon lines

diminish rapidly, and for $\lambda > 4.5 \mu\text{m}$ there are no suitably strong argon lines that can be used for wavelength calibration. Therefore, to calibrate the LXD modes, we have adopted an approach that makes use of both the argon lines at shorter wavelengths and the numerous sky emission lines at longer wavelengths. A pure sky frame (for A-sky mode) or two target frames (for AB mode) are used to produce a night-sky emission image. If the object is observed in the AB mode, then this is done by combining the frames according to the equation

$$\text{sky} = (\mathbf{A} + \mathbf{B}) - |(\mathbf{A} - \mathbf{B})|. \quad (5)$$

For those LXD orders (4 and 6) where the arc lines are either

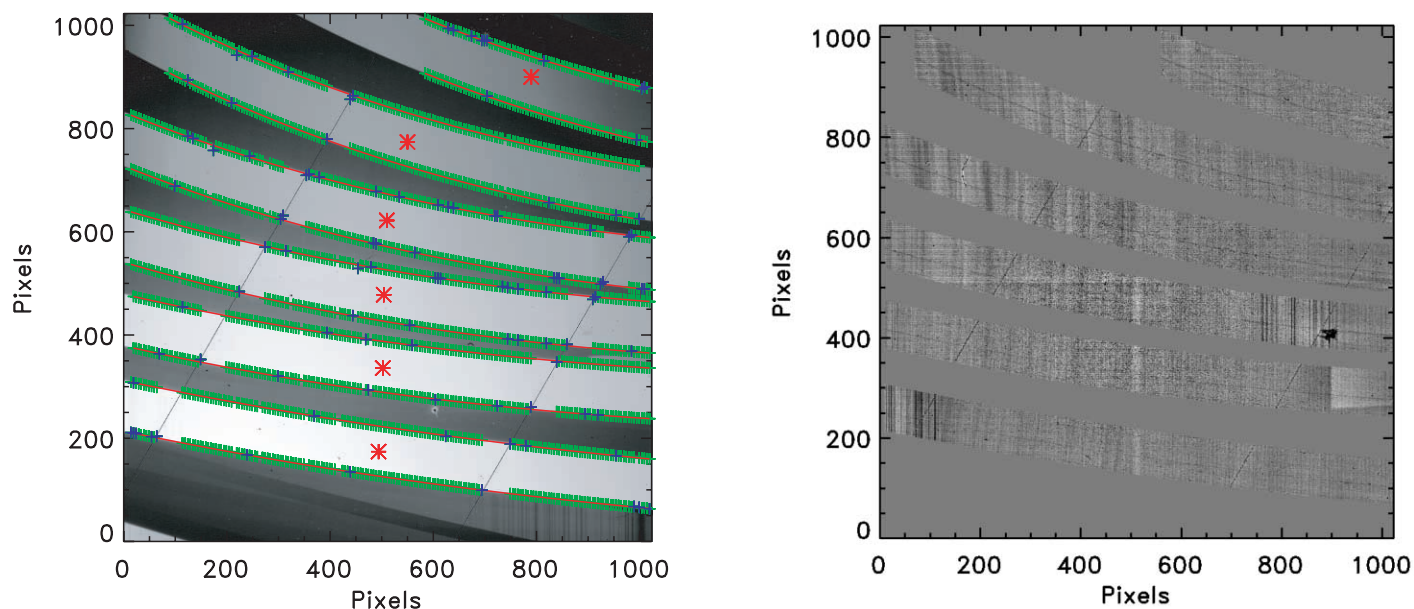


FIG. 4.—*Left*: Detection of the edges of the orders in the flat-field image. The red asterisks represent initial guesses for the centers of each order. The green crosses mark the positions where the order edge is located, and the red lines are the robust polynomial fits through these positions. The blue crosses designate positions that have been excluded from the fits. *Right*: A normalized flat field. Note the tree-ring-like structure caused by mechanical thinning of the InSb array material from 200 to 20 μm , and the odd-even eight column pattern due to the eight separate readout locations in the silicon multiplexer. Also visible is a slight difference in gain between the top half and bottom half of the array, which are read out by two different analog-to-digital converter boards. The spot at the approximate coordinates (900, 400) is a clusters of bad pixels.

weak or absent, the sky emission image replaces the arc image. The final image is then written to disk. Figure 5 shows an arc image for LXD2.3 mode.




FIG. 5.—LXD2.3 wavelength calibration frame. The five orders (4–8; *bottom to top*) are apparent. Orders 4 and 6 have been replaced by a sky emission spectrum.

4.2. Step 2: Spectral Extraction

4.2.1. Image Processing

The first step in processing the science images involves correcting each pixel value in the raw data frames for possible nonlinearity. Spextool accomplishes this using a nonlinearity curve constructed by P. Onaka and J. Leong from flat-field data taken expressly for this purpose. A series of flats was taken with gradually increasing exposure times, and the nonlinearity curve was then determined by fitting a third-order polynomial to the observed count rate as a function of the observed counts in each pixel. The deviation of this curve from a constant value then gives the factor needed to correct the observed counts in a given pixel to the expected counts for a linear response.

Unfortunately, as noted by Vacca et al. (2004), the linearity correction factor cannot be applied directly to the stored pixel values, because the array is read out using multiple correlated double sampling (MCDS; Fowler & Gatley 1990), also known as Fowler sampling. In this readout mode, the array is read out nondestructively n_r times at the start of the integration. After a time Δt (the integration time), the array is again read out

nondestructively n_r times. Each nondestructive read requires a time of δt . The *net signal*, given by the difference between the sums of the values recorded for the two sets of nondestructive reads, is then written to disk as an image. If the user requests n_c “co-adds,” then the readout process is repeated n_c times, and the individual net signals are summed in memory and the total net signal \mathbf{S}_{tot} is written to disk. The linearity correction should be applied to the values of the individual nondestructive reads, but only the total net signal is written to disk as an image. Therefore, Spextool determines the nonlinearity correction for each pixel in each input frame using the iterative technique described by Vacca et al. (2004).

The linearity-corrected data $\mathbf{S}_{\text{tot}}^c$ are then converted to a “flux” \mathbf{I} (DN s⁻¹) given by

$$\mathbf{I} = \frac{\mathbf{S}_{\text{tot}}^c}{n_r n_c \Delta t}. \quad (6)$$

Because optimal extraction (see § 4.2.5) requires an error estimate for each pixel in the extraction aperture, Spextool also constructs a variance image \mathbf{V}_i for each input image, given by

$$\mathbf{V}_i = \frac{\mathbf{S}_{\text{tot}}^c}{g n_r n_c^2 \Delta t^2} \left[1 - \frac{1}{3} \frac{\delta t}{\Delta t} \frac{(n_r^2 - 1)}{n_r} \right] + \frac{2\sigma_{\text{read}}^2}{g^2 n_r n_c \Delta t^2} \quad (7)$$

where g is the gain (e^- DN⁻¹) and σ_{read} is the rms read noise (e^-). A derivation of equation (7) is given by Vacca et al. (2004).

The next step consists of subtracting an \mathbf{I}_A and \mathbf{I}_B frame to remove the additive components of the recorded signal as described in § 3. The resulting image is then flat-fielded using the normalized flat field created in § 4.1.1. This yields the final reduced image

$$\mathbf{D} = \frac{\mathbf{I}_A - \mathbf{I}_B}{\mathbf{flat}}. \quad (8)$$

The variance image of \mathbf{D} is given by,

$$\mathbf{V}_D = \frac{\mathbf{V}_A + \mathbf{V}_B}{\mathbf{flat}^2}. \quad (9)$$

A pair-subtracted, flat-fielded image of a point source in the SXD mode is shown in Figure 6. Note the positive and negative spectra produced by the pair subtraction.

4.2.2. Spatial Profiles

The spatial profile image $P_{i,j}$ gives the relative intensity of a target spectrum at each spatial position j along the slit at each column (wavelength) i . The spatial profile varies as a function of both wavelength (column) and spatial position (row) because of small optical effects, atmospheric differential refraction, and wavelength-dependent seeing, which cause changes in the rel-




FIG. 6.—Pair-subtracted, flat-fielded SXD frame. The source spectra in each order are the strong positive and negative curved lines. Note the regions of zero source flux due to strong atmospheric absorption. The faint horizontal lines seen at columns above 900 in rows 300 and 350 are K -band “ghosts” due to internal reflections in SpeX.

ative shape and position of the spatial profile (Horne 1986). Constructing an accurate, high-S/N spatial profile image is crucial for the optimal extraction procedure. In addition, the mean spatial profile along the slit (averaged over all columns i) is needed by the user to define the extraction parameters (see § 4.2.4). The procedure used by Spextool to construct the spatial profile images is a variation on the “virtual resampling” algorithm of Mukai (1990).

The spatial profile image for each order is constructed as follows. If the target is a point source, then the median signal level in each column (within an order) is subtracted from each pixel in the column to remove any residual sky background. If the target is an extended source, then no attempt is made to subtract any residual sky signal, because the sky signal cannot be measured if the target subtends an angle greater than 15”, the length of the slit. The order is then resampled onto a uniform rectangular grid of i columns and s spatial points, as described in § 4.1.1. Figure 7 (*top left*) shows the resampled image of order 6 in the SXD mode.

The resampled image of an order can be thought of as a spatial profile image with an intensity modulated by the observed spectrum of the object at each wavelength. Therefore, at any given column i and spatial point s in the order image, the intensity is given by

$$O_{i,s} = f_i P_{i,s}, \quad (10)$$

where f_i is the intensity of the source spectrum at column (wave-




FIG. 7.—*Top left*: Resampled image of order 6 of the SXD mode. Note the positive and negative spectrum produced by the pair subtraction in the AB mode. The sharp peaks are bad pixels. *Top right*: First-order approximation of the spatial profile image. *Bottom left*: Model of the spatial profile image created by evaluating the profile coefficients. *Bottom right*: Average spatial profile.

length) i . If f_i is known, then $P_{i,s}$ can be generated easily by simple division. (Of course f_i is exactly what we are trying to determine from the observations and so is never known a priori!) The variation in intensity with column can be removed by normalizing each column in $O_{i,s}$ by an appropriate scale factor. For optical spectroscopy, the scale factor is given by the total intensity in a column, which can be calculated by simply summing the values of the pixel along the column. Dividing each pixel in a column by the total flux in that column removes the variation in wavelength, and yields $P_{i,s}$ (Horne 1986). However, in the near-infrared, the AB mode of observing produces both positive and negative images of the spectrum in any column when the two exposures are subtracted, and as a result, $\sum_s f_{i,s} \approx 0$ for all i . Therefore, Spextool computes the profile image in the following iterative manner.

First, at each spatial (slit) point s , Spextool computes the

median of the columns of the observed order image. This yields a one-dimensional array (of length s) that is the first-order approximation to the mean spatial profile P_s . Assuming that

$$O_{i,s} \approx c_i P_s, \quad (11)$$

Spextool then uses a linear least-squares routine to fit the first-order approximation of P_s to the intensity variation in each column in the resampled order image and determines the scale factors c_i . These scale factors are actually the first-order approximations to f_i . The resampled image is then divided by the scale factors, which then yields an approximation to $P_{i,s}$. The result of this procedure, applied to order 6, is shown in Figure 7 (*top right*). The entire procedure can be repeated to yield better approximations to $P_{i,s}$. However, we have found that this

first $P_{i,s}$ image is quite satisfactory for optimal extraction purposes. Further iterations produce only marginal improvements.

Once the spatial profile image $P_{i,s}$ is determined, Spextool fits a low-order polynomial along the columns i at each spatial point s in the image. This has the effect of smoothing the profile image, thus increasing its S/N. Columns for which the scale factors c_i indicate there is no significant source flux are excluded from the polynomial fit. The *profile coefficients* resulting from the polynomial fits can be used to determine the value of the intensity profile $P_{i,s}$ at any column and any spatial point. The spatial profile image for order 6, created by evaluating the profile coefficients, is shown in Figure 7 (*bottom left*). The average spatial profile P_s , computed from the model spatial profile image, is shown in Figure 7 (*bottom right*). Finally, $P_{i,j}$ is constructed by resampling $P_{i,s}$ onto the pixel grid (i, j) using the arcsecond values associated with each pixel.

4.2.3. Aperture Location and Tracing

Once the profile coefficients have been determined, the user specifies the position of the extraction apertures along the slit using the average spatial profiles as a guide. If the target is a point source, Spextool automatically determines the aperture positions by iteratively identifying local maxima in the absolute value of the average spatial profiles. If the target is an extended source, the user types the locations of the aperture positions, measured in arcseconds from the bottom of the slit (see Fig. 8), into a text field.

Since the location of the extraction apertures are identified using the average spatial profiles, they do not represent the true position of the spectrum on the detector at any given column. As discussed in § 4.2.2, this is a result of small optical effects as well as atmospheric differential refraction. For point sources, this true position is found by tracing the spectra across the array. A tracing routine steps across the array at regularly spaced intervals and locates the center of a spectrum by fitting a Gaussian in the spatial dimension centered on the aperture position defined above. The individual center points are then fitted as a function of column number with a low-order robust polynomial. The resulting *trace coefficients* provide the location of that spectrum at any column. This process is repeated for each aperture and for each order.

Unfortunately, extended sources cannot, in general, be traced, since there is no guarantee that the extraction apertures will be located at the positions of local maxima in the average spatial profiles. Therefore, Spextool simply adopts trace coefficients based on the edge coefficients. Spextool first converts the user-specified aperture position (in arcseconds) to a pixel position at each column using the edge coefficients. The aperture positions (in pixels) are then fitted as a function of column number with a low-order robust polynomial to produce the trace coefficients. Therefore, due primarily to atmospheric differential refraction, the position of the extraction aperture on the target will vary with wavelength. For most air masses,

this effect can be mitigated by choosing a large enough extraction aperture.

4.2.4. Extraction Parameters

The extraction parameters (i.e., extraction apertures and background regions) are defined by the user with the aid of the average spatial profiles constructed in § 4.2.2. Figure 8 shows the average spatial profiles for a point source and an extended source (Uranus) in order 6 of the SXD mode.

For point sources, the parameters are defined with respect to each aperture position (*vertical cyan solid line*). The following extraction parameters are indicated in color in Figure 8:

PSF radius.—The radius in arcseconds at which the PSF profile falls to zero (*blue dotted lines*). This parameter is used only for optimal extraction (see § 4.2.5).

Aperture radius.—The radius in arcseconds defining the edge of the extraction aperture (*green solid and vertical dotted lines*).

Background start.—The radius in arcseconds at which to start the background region.

Background width.—The width in arcseconds of the background region (the background region is shown with a red solid line).

Background fit degree.—The degree of the robust least-squares polynomial fit to the background regions (the fit to the background is shown with a magenta dotted line).

For extended sources, the relevant parameters are:

Aperture radii.—The extraction radius for each aperture (*green solid and dotted line*).

Background regions.—The regions used to fit the background level (*red solid lines*).

Background fit degree.—The degree of the robust least-squares polynomial fit to the background regions (*magenta dotted line*).

4.2.5. Extracting the Spectra

The spectra can be extracted once the trace coefficients are determined and the extraction parameters have been defined by the user. For a given aperture and order, the trace coefficients are used to determine the position of the spectrum at a given column. The extraction parameters, defined in arcseconds, are then converted to units of pixels using the edge coefficients. Before the spectra can be extracted, any residual sky signal left over from the initial pair subtraction (see § 3) must be subtracted from the reduced image \mathbf{D} (eq. [8]). At each column i , the values of the pixels in the user-defined background regions (red solid lines in Fig. 8) are fitted with a low-order weighted polynomial as a function of their position along the slit. The resulting fit provides an estimate of the residual sky level $\tilde{R}_{i,j}$ at each pixel along the slit in column i . As long as the raw pixel values in the individual raw \mathbf{A} and \mathbf{B} frames are




FIG. 8.—Two examples of the average spatial profiles for order 6 of the SXD mode. The upper profile results from a point source observed in the AB mode; the bottom is an extended source (Uranus) observed in the A-sky mode. The solid vertical cyan lines denote the location of the two apertures. Spextool automatically identified the positions of the two apertures in the upper panel, while we manually selected two apertures located at 6"3 and 9"5 in the lower panel. The solid and dotted green lines denote the extraction aperture, while the red lines denote the regions used for the background determination. The magenta dotted line represents the polynomial fit to the background regions. The dotted blue lines in the point-source profile represent the PSF radius. The PSF radius is used only for optimal extraction and is therefore not present in the extended source profile.

not saturated, subtracting $\tilde{R}_{i,j}$ from $D_{i,j}$ removes the residual sky signal.

Spextool can extract the spectra from the sky-subtracted frame ($D_{i,j} - \tilde{R}_{i,j}$) in two ways: *sum* and *optimal*. In the sum mode, cosmic-ray hits are first identified using the spatial profiles created in § 4.2.2. At each column i , the spatial profile $P_{i,j}$ is scaled to the data in a robust linear least-squares sense. Any pixels identified as outliers, as well as bad pixels identified from the bad pixel mask stored on disk, are replaced with the value of the scaled profile. The flux and variance of the object at column i is then given by

$$f_i^{\text{sum}} = \sum_{j=j_1}^{j_2} \epsilon_{i,j} (D_{i,j} - \tilde{R}_{i,j}),$$

$$V_{f_i^{\text{sum}}} = \sum_{j=j_1}^{j_2} \epsilon_{i,j}^2 (V_{D_{i,j}} + V_{\tilde{R}_{i,j}}), \quad (12)$$

where j_1 and j_2 are the pixel indices of the upper and lower limits of the extraction aperture, and $\epsilon_{i,j}$ is the fractional weight of pixel (i,j) , which is given by the fraction the each pixel

contained within the extraction aperture. The value of $\epsilon_{i,j}$ is everywhere unity except for the pixels at the edges of the apertures.

Optimal extraction was originally developed by Robertson (1986) and Horne (1986) to maximize the S/N of long-slit spectra of point sources. The method was later expanded by Marsh (1989) and Mukai (1990) to include the extraction of cross-dispersed spectra, and Naylor (1998) has developed an “optimal extraction” algorithm for point-source photometry. All optimal extraction algorithms work on the same principle, namely, weighting each pixel in the extraction aperture in proportion to the amount of flux it contains from the target.

For each aperture, Spextool normalizes $P_{i,j}$ using the PSF radius defined in § 4.2.4 so that

$$P_{i,j} = \frac{P_{i,j}}{\sum_{j=j_1}^{j_2} P_{i,j}}, \quad (13)$$

where j_1 and j_2 are the upper and lower limits given by the PSF radius. Here $P_{i,j}$ now gives the fraction of the total flux from the object contained in pixel j of column i . At each column

i , $(D_{i,j} - \tilde{R}_{i,j})/P_{i,j}$ then represents an array of j independent estimates of the total flux of the object. These independent estimates are combined in an optimal way using a weighted mean where the weights are given by their inverse variances, $P_{i,j}^2/(V_{D_{i,j}} + V_{\tilde{R}_{i,j}})$. The flux and variance of the object at column i then become

$$f_i^{\text{opt}} = \frac{\sum_{j_1}^{j_2} M_{i,j} P_{i,j} (D_{i,j} - \tilde{R}_{i,j}) / (V_{D_{i,j}} + V_{\tilde{R}_{i,j}})}{\sum_{j_1}^{j_2} M_{i,j} P_{i,j}^2 / (V_{D_{i,j}} + V_{\tilde{R}_{i,j}})},$$

$$V_{f_i^{\text{opt}}} = \frac{1}{\sum_{j_1}^{j_2} M_{i,j} P_{i,j}^2 / (V_{D_{i,j}} + V_{\tilde{R}_{i,j}})}, \quad (14)$$

where j_1 and j_2 are the upper and lower limits given by the aperture radius, and $M_{i,j}$ is an optional mask to remove bad pixels.

As explained by Horne (1986), optimal extraction provides an increase in the S/N only for background-limited spectroscopy. Since the SXD mode is source limited for all but the faintest targets, optimal extraction provides relatively little increase in the S/N of the spectra. However, the LXD mode is almost always background limited because of the strong thermal flux of the sky and telescope, and as a result, optimal extraction provides an increase in the S/N of the spectra of up to 25%. Independent of S/N concerns, optimal extraction is superior to sum extraction because it ignores bad pixels instead of trying to fix them.

4.2.6. Wavelength Calibration

The spectra are wavelength calibrated using the arc frame created in § 4.1.2. As discussed there, we use argon lines and sky emission features for wavelength calibration. We have compiled a list of argon lines from Rao, Humphreys, & Rank (1966) and have removed lines that are blended at $R = 2500$, the maximum resolving power available with SpeX. We determined the wavelengths of the sky emission features using a model of the atmospheric transmission computed with ATRAN (Lord 1992). All wavelengths are given in vacuum.

Spextool begins the wavelength-calibration process by extracting arc spectra for each extraction aperture. Argon lines or sky features that are blended at the resolving power of the object spectra are then pruned from the line list. The approximate position of each line in the pruned line list is determined from a previously determined wavelength solution stored on disk. Due to the combined effects of the nonrepeatability of the slit wheel mechanism, flexure in the instrument, and internal movements due to the thermal recycling of SpeX, the zero point of this solution can vary by up to a few pixels. Spextool, therefore, performs a cross-correlation between the arc spectrum extracted above and an arc spectrum stored on disk (used to create the solution stored on disk) to determine the zero point. The position of each line is then determined by fitting

a Gaussian to a small section of the arc spectrum centered on the approximate position of each line.

Spextool then determines a single wavelength solution for each extraction aperture using a variation of the technique outlined by Hall et al. (1994). This method has two advantages over determining a wavelength solution on an order-by-order basis. First, the overall fit is of higher quality, because all of the lines on the array are used in the fit. Second, orders with a small number of lines can still be accurately calibrated. The method makes use of the fact that for a fixed groove spacing, the ratio of two wavelengths λ_1 and λ_2 , diffracted at the same angle but located in two different orders m_1 and m_2 , is given by

$$\frac{\lambda_1}{\lambda_2} = \frac{m_2}{m_1}. \quad (15)$$

Since we assume the spatial dimension is aligned with the columns of the detector (see § 2), a given column corresponds to the same diffraction angle for all of the orders. Therefore, Spextool scales the wavelengths of the identified lines to their equivalent wavelengths λ' in a given order m' using equation (15). For example, in the SXD mode, a line with $\lambda = 2.2 \mu\text{m}$ in order 3 will have $\lambda' = 1.1 \mu\text{m}$ for $m' = 6$.

In principle, a wavelength solution can be computed by fitting the scaled lines (λ') as a function of column number and then applying the solution to the individual orders using equation (15). However, we find that such a fit leaves large-scale trends with order number in the residuals (due to small optical effects). To remove these trends, we fit the scaled wavelengths with a robust two-dimensional polynomial where λ' is a function of both the column number i and the order number m :

$$\lambda'(i, m) = \sum_{k=0}^{N_k} \sum_{l=0}^{N_l} C_{kl} i^k m^l. \quad (16)$$

A third-order polynomial in both i and m ($N_k = N_l = 3$) removes the trends found in the residuals of the one-dimensional fit. Figure 9 shows the residuals of the wavelength calibration fit for the SXD mode. The rms of the residuals is 0.075 \AA . The rms of the residuals for the LXD1.9 modes is roughly an order of magnitude greater, owing to the fewer lines available for wavelength calibration.

For slit widths wider than $0''.8$, the argon and sky lines become hopelessly blended, which results in a poor fit, due to the small number of individual lines that can be identified. In this case, the user can adopt a wavelength solution stored on disk after applying the cross-correlation zero point.

4.2.7. Results

The extracted spectrum from order m is then wavelength calibrated by scaling the results of evaluating equation (16) by m'/m . The resulting wavelength-calibrated spectra (including




FIG. 9.—Residuals of the wavelength-calibration procedure for the SXD mode, shown as a function of spectral order (*top*) and column (*bottom*). The rms of the residuals is 0.075 Å.

the error spectra) are written to disk as a FITS file. All of the pertinent extraction parameters are written to the FITS header. Figure 10 shows the SXD and LXD2.1 spectra of the B5 sub-giant HD 147394.

4.3. Step 3: Post-Extraction Processing

4.3.1. Combining the Spectra

For each input image, Spextool produces an output file containing the wavelength-calibrated flux and error spectra corresponding to each extraction aperture in each order. Multiple output files (resulting from repeated observations) can be combined with a routine that allows the user to choose the combining procedure (mean, weighted mean, or median). The spectra are resampled onto a common wavelength grid and then combined.

4.3.2. Telluric Correction

In contrast to the relatively weak telluric absorption features present in the optical, strong absorption bands due to H₂O, CO₂, and CH₄, and to a lesser extent N₂O, O₂, and O₃, dominate the 1–5 μm wavelength range. Therefore, removing the telluric absorption features in the spectrum of the target is both a necessary and critical step in the reduction of infrared spectra.

We have developed a new method for correcting medium-resolution, near-infrared spectra that makes use of an A0 V star, observed near in time and close in air mass to the target, and a high-resolution model of Vega. This method simultaneously removes the instrument signature and approximately flux calibrates the target spectrum. A full explanation of the method, as well as examples of the process, is given by Vacca, Cushing, & Rayner (2003).

4.3.3. Merging the Orders

Once the target spectra have been corrected for telluric absorption, the spectra from the various orders can be merged into a single, continuous spectrum. In principle, the flux levels and slopes of the spectra from two neighboring orders should match exactly in their common wavelength range. However, we find significant differences in both the flux levels and slopes for some objects when the nonlinearity correction is not incorporated into the reduction process (§ 4.2.1). For example, the top panel of Figure 11 shows the mismatch in flux and slope between the uncorrected spectra from orders 6 (*green*) and 7 (*red*) of Gl 846 (M0.5 V). The nonlinearity correction removes this problem to a large degree, as can be seen in the lower panel of Figure 11. For most objects the corrected flux levels in the




FIG. 10.—*Left*: SXD spectra of HD 147394 (B5 IV) prior to telluric correction. The overall shape of the spectra is due to the instrument throughput. Strong H absorption lines can be seen in orders 6 and 7. The high-frequency structure is due to telluric absorption. *Right*: LXD2.1 spectra of HD 147394 prior to telluric correction. Most of the structure in these spectra is due to telluric absorption.

order overlap regions differ by $\leq 2\%$. To account for this small offset before the spectra from two orders are merged, the user can scale one of the spectra to match the flux level of the other. One spectrum is then resampled onto the wavelength grid of the other, and the two spectra are combined using a weighted mean. The process is continued until all of the orders have been merged.

4.3.4. Cleaning and Smoothing the Spectrum

Cleaning and smoothing the spectra are cosmetic in nature and are therefore optional. In principle, the extraction process either ignores or fixes bad pixels. However, occasionally bad pixels are missed and corrupt the spectrum. In addition, the division by the telluric correction spectrum can produce strong, high-frequency noise deep in the telluric bands, since the flux values in these wavelength ranges can be close to zero. Therefore, the user can either remove sections of the spectrum that are noisy, or interpolate over single bad pixels.

For faint objects, it is often advantageous to smooth or rebin a spectrum to increase its S/N. The user has the option of smoothing the spectrum using a Gaussian kernel or a Savitzky-Golay kernel (Press et al. 1992). The advantage of the latter is that it increases the S/N of the spectrum with minimal loss of resolution.

4.3.5. Merge SXD and LXD Spectra

Finally, if the object has been observed in both the SXD and LXD mode, the spectra from the two modes can be combined to produce a single, continuous spectrum. In principle, the flux levels of the SXD and LXD spectra should match, but in practice we often find that the absolute flux levels differ. Therefore, before the spectra from the two modes are combined, one spectrum can be scaled to match the flux level of the other. The LXD spectrum is then resampled onto the wavelength grid of the SXD spectrum, and the two spectra are combined using a weighted mean.




FIG. 11.—Spectra of Gl 846 (M0.5 V) from orders 6 and 7 of the SXD mode. The spectra in the upper panel were reduced without applying a correction for nonlinearity; the corrections were incorporated into the reductions of the spectra in the bottom panel. Incorporation of the nonlinearity corrections removes, to a large degree, the mismatch in both the flux levels and slopes of the spectra in the overlap regions. No additional scaling has been applied to these spectra.

4.3.6. Results

Figure 12 shows the approximately flux-calibrated 0.8–5.0 μm spectrum of HD 147394 (B5 IV). The resolving power of the 0.8–2.5 μm spectrum is $\langle R \rangle = 2000$, and the S/N is greater than 100; for the 2.1–5.0 μm spectrum, $\langle R \rangle = 2500$, and the S/N is ~ 70 . The break in the spectrum centered at 1.83 μm is a result of a gap in the wavelength coverage of the SXD mode (see Rayner et al. 2003). The spectrum from 2.5 to 2.9 μm and from 4.2 to 4.5 μm has been removed because the atmosphere is almost completely opaque at these wavelengths. The Brackett, Paschen, and Pfund series have been identified.

5. SUMMARY

We have presented a description of a new IDL package called Spextool, which has been developed for the reduction of spectral data taken with the near-infrared, medium-resolution, cross-dispersed spectrograph SpeX at the IRTF. Spextool allows rapid reduction (i.e., flat fielding, aperture definition, sky subtraction, spectral extraction, wavelength calibration, and various postprocessing steps) of spectral data acquired in the most commonly used spectroscopic modes of SpeX. It generates realistic estimates of error arrays associated with spectra, and incorporates an “optimal extraction” algorithm for point-source




FIG. 12.—The 0.8–5.0 μm spectrum of HD 147394 (B5 IV). The break in the spectrum centered at 1.83 μm is a result of a gap in the wavelength coverage of the SXD mode. The spectrum from 2.5 to 2.9 μm and from 4.2 to 4.5 μm has been removed because the atmosphere is almost completely opaque at these wavelengths. The hydrogen absorption lines from the Paschen, Brackett, and Pfund series are indicated. Note the excellent agreement in flux between the separate spectral regions.

data. A user interacts with Spextool via graphical user interfaces that incorporate a set of buttons and alphanumeric text fields that specify the various parameters needed for the spectral extraction procedures.³

³ The most current version of the package (ver. 3.2) can be downloaded from the SpeX website, <http://irtfweb.ifa.hawaii.edu/Facility/spex>. In addition,

We thank Marc Buie, Nick Kaiser, and John Tonry for helpful discussions, and Peter Onaka and Jonathan Leong for their work on the nonlinearity curve. M. Cushing acknowledges financial support from the NASA Infrared Telescope facility. W. Vacca thanks Alan Tokunaga and James Graham for support.

detailed instructions on how to run Spextool are included in the package and can also be downloaded from the SpeX website.

REFERENCES

- Fowler, A. M., & Gatley, I. 1990, *ApJ*, 353, L33
 Hall, J. C., Fulton, E. E., Huenemoerder, D. P., Welty, A. D., & Neff, J. E. 1994, *PASP*, 106, 315
 Horne, K. 1986, *PASP*, 98, 609
 Joyce, R. R. 1992, in *ASP Conf. Ser. 23, Astronomical CCD Observing and Reduction Techniques*, ed. S. B. Howell (San Francisco: ASP), 258
 Lauer, T. R. 1999, *PASP*, 111, 1434
 Lord, S. D. 1992, *A New Software Tool for Computing Earth's Atmospheric Transmission of Near- and Far-Infrared Radiation* (NASA Tech. Mem. 103957; Washington: NASA)
 Marsh, T. R. 1989, *PASP*, 101, 1032
 Mukai, K. 1990, *PASP*, 102, 183
 Naylor, T. 1998, *MNRAS*, 296, 339
 Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, *Numerical Recipes in Fortran 77* (New York: Cambridge Univ. Press)
 Rao, K. N., Humphreys, C. J., & Rank, D. H. 1966, *Wavelength Standards in the Infrared* (New York: Academic Press)
 Rayner, J. T., Toomey, D. W., Onaka, P. M., Denault, A. J., Stahlberger, W. E., Vacca, W. D., Cushing, M. C., & Wang, S. 2003, *PASP*, 115, 362
 Robertson, J. G. 1986, *PASP*, 98, 1220
 Schroeder, D. J. 2000, *Astronomical Optics* (San Diego: Academic Press)
 Vacca, W. D., Cushing, M. C., & Rayner, J. T. 2003, *PASP*, 115, 389
 Vacca, W. D., Cushing, M. C., & Rayner, J. T. 2004, *PASP*, 116, 352