The Astrophysical Journal Supplement Series, 147:379–401, 2003 August
© 2003. The American Astronomical Society. All rights reserved. Printed in U.S.A.

A UNIFORM DATABASE OF 2.4–45.4 MICRON SPECTRA FROM THE *INFRARED SPACE OBSERVATORY* SHORT WAVELENGTH SPECTROMETER[1]

G. C. Sloan,[2,3] Kathleen E. Kraemer,[4,5] Stephan D. Price,[4] and Russell F. Shipman[6]

Received 2002 June 19; accepted 2003 March 6

ABSTRACT

We present a complete set of all valid SWS full-scan 2.4–45.4 μm spectra processed and renormalized in as uniform a manner as possible. The processing produces a single spectrum for each observation from the 288 individual spectral segments, which are the most processed form available from the *ISO* archive. The spectra, and the programs used to create them, are available to the community on-line.

Subject headings: atlases — infrared: general — methods: data analysis — techniques: spectroscopic

1. INTRODUCTION

The data archive for the *Infrared Space Observatory* (*ISO*) contains 1271 observations taken with the Short Wavelength Spectrometer (SWS) in full-scan mode, which produced moderate-resolution full-grating spectra from 2.38 to 45.4 μm. These data are available from the *ISO* Data Archive (IDA) in partially processed form, as are interactive tools for further processing. Spectra reduced with these tools have led to excellent scientific results and a multitude of papers. However, the interactivity of the *ISO* data reduction software, such as the Observer's SWS Interactive Analysis (OSIA) system and the *ISO* Spectral Analysis Package (ISAP), precludes the reprocessing of large numbers of spectra. Also, one person starting with a single observation might not produce the same final spectrum as another starting with the same data, due to the inherent subjectivity of some of the processing steps. Custom-designed software is usually optimized for a single spectrum or a small number of spectra and is generally not available for others to apply to their own data.

The STARTYPE program developed as a series of proposals to obtain spectra, primarily with the SWS, on *ISO* to achieve two goals: (1) to obtain spectra of sources which represent as many different spectral morphologies as possible and develop an infrared spectral classification system from those observations, and (2) to make available to the astronomy community an atlas of uniformly processed SWS spectra. Kraemer et al. (2002, hereafter Paper I) present the infrared spectral classifications of nearly all of the SWS full-scan spectra. This paper describes the production of the spectral atlas and explains how the community can access any of the spectra within it.

These two projects evolved together, with the early results of the spectral processing influencing the classifications and the final classifications influencing how particular groups of spectra were processed. The two works complement each other in another way, as the classifications facilitate the identification of groups of sources showing similar spectral properties which can then be studied using spectra obtained from the atlas. The atlas presents each spectrum in a more completely processed format than was previously available through the IDA.

Our examination of the IDA finds that the results of analysis from only ~450 of the more than 1250 full-scan SWS spectra have been published (as of 2002 November), which leaves nearly two-thirds of the database untreated in the literature. The IDA provides the means for authors of previously published spectra to make their processed data available to the community, but most authors have not yet done so. While it is unlikely that the automated processing described here will provide results superior to individual spectra carefully processed interactively, the automated processing does make the entire database available in a user-friendly format for the first time. The atlas is particularly suited to population studies, especially when used in conjunction with the classifications in Paper I.

The recent study of Group 2.SE (stars embedded within optically thin oxygen-rich dust shells) by Sloan et al. (2003) provides one example of this approach. Paper I presented a two-level classification system, where the first level depended on the overall shape of the spectral energy distribution (SED) and the second level identified the dominant spectral component. Paper I also explained how more detailed analysis would lead to a third level of classification where the spectral differences among the members of a given class might be examined and perhaps ordered into a sequence. For 2.SE sources, Sloan & Price (1995, 1998) have already performed this level 3 classification (using *IRAS* data), so Sloan et al. (2003) have focussed on spectral properties associated with the 13 μm dust emission feature within that context. Heras et al. (2002) provide a similar analysis for 1.N and 1.NO sources (naked stars with atomic lines and oxygen-rich molecular absorption bands in their spectra, respectively). Much of their analysis was complete before our database was ready, but they did supplement their sample using the classifications in Paper I, and they have made suggestions on how the level 3 classification of

[1] Based on observations with the *Infrared Space Observatory*, a European Space Agency (ESA) project with instruments funded by ESA Member States (especially the Principal Investigator countries: France, Germany, the Netherlands, and the United Kingdom) and with the participation of the Institute of Space and Astronautical Science (ISAS) and the National Aeronautics and Space Administration (NASA).

[2] Institute for Scientific Research, Boston College, Chestnut Hill, MA 02467.

[3] Infrared Spectrograph Science Center, Cornell University, Ithaca, NY 14853-6801; sloan@isc.astro.cornell.edu.

[4] Air Force Research Laboratory, Space Vehicles Directorate, 29 Randolph Road, Hanscom AFB, MA 01731; kathleen.kraemer@hanscom.af.mil, steve.price@hanscom.af.mil.

[5] Institute for Astrophysical Research, Boston University, Boston, MA 02215.

[6] SRON, Space Research Institute of the Netherlands, Groningen, the Netherlands; russ@sron.rug.nl.

these oxygen-rich naked stars should proceed. This database will help facilitate the follow-up. One can envision similar studies of sources embedded within optically thin carbon-rich dust shells (2.CE) or optically thick carbon-rich shells (3.CE and 3.CR), or other classes of objects such as planetary nebulae showing the infrared emission features attributed to polycyclic aromatic hydrocarbons (4.PU).

Section 2 describes the SWS, the format of the data from the IDA, and how our algorithm reprocesses these data. Section 3 describes the resulting spectral atlas. Section 4 validates the results, while § 5 discusses limitations and caveats to the database. The appendices explain how to access the on-line atlas and describe the algorithm in more detail.

2. THE POST–PIPELINE PROCESSING ALGORITHM

2.1. *The Short Wavelength Spectrometer and the Pipeline*

The SWS is a scanning spectrometer. Leech et al. (2002) provide a thorough description of the SWS. We briefly review the points relevant to the post–pipeline processing algorithm described below. Of the several Astronomical Observing Templates (AOTs), the SWS observed 1248 of the spectra using AOT 1, which scanned wavelengths from 2.35 to 45.4 μm onto four arrays of 12 detectors each. AOT 1 used these arrays to create 12 discrete spectral segments (or bands), varying in length from 0.25 μm for Band 1A (2.35–2.60 μm) to nearly 18 μm for Band 4 (27.7–45.4 μm; see Table 1). Each segment consists of data from 12 detectors, scanned in both positive and negative wavelength directions, giving 24 individual spectra for each of the 12 segments. Thus, a single observation consists of 288 spectra which must be combined into a single coherent spectrum.

An additional 23 full-scan spectra were observed using AOT 6, a more sensitive mode with higher resolution usually devoted to scanning smaller wavelength ranges. In its full-scan mode, it operated very similarly to AOT 1. Appendix C explains how the processing of the AOT 6 data differed from the AOT 1 data.

The IDA provides spectra from the SWS in the form of Auto-Analysis Results (AARs) produced by the Off-Line Processing (OLP) Pipeline. Our data reduction algorithm begins with output from OLP version 10.1, which was released in 2001 November. The pipeline calibrates the raw data, converts them into astronomical units (flux density in Jy, wavelength in μm), corrects for a number of instrumental effects, and writes the results as AAR files. Each AAR file contains the data from one observation (known as a TDT for target dedicated time) as a series of wavelength elements. For each wavelength element, the AAR gives wavelength, flux, uncertainty, a series of error flags, the detector number, scan direction, scan number, and segment number.

Leech et al. (2002, Table 2.1), give a range of resolutions for the SWS between 930 and 2450 ($\lambda/\Delta\lambda$). The resolution changes significantly from one segment to the next. The values quoted are appropriate to an extended source observed in the highest resolution possible in AOT 1, speed 4. Speed 3 halves the resolution, and speeds 1 and 2 cut the resolution in half again. Speeds 1 and 2 provide the same resolution, but speed 2 was scanned more slowly to achieve greater sensitivity.

Figure 1 presents typical pipeline results for two sample spectra and illustrates two immediately apparent problems. First, the data in each spectral segment show significant scatter, due primarily to the difficulty of properly aligning spectra from different detectors and scan directions. The scatter from a single detector in one scan direction is significantly smaller. Second, discontinuities appear between spectral segments. Our post–pipeline processing algorithm, which we will refer to as the *swsmake* code, handles these two problems separately. First, it combines the 24 spectra from all 12 detectors and both scan directions in each spectral segment into one. Then it combines the 12 spectral segments into a single continuous spectrum by normalizing the segments to each other.

TABLE 1
THE SWS SPECTRAL SEGMENTS

Segment	Band	Detector	λ Range (μm)	Aperture (arcsec)
1	1A	InSb	2.35–2.60	14 × 20
2	1B	InSb	2.60–3.02	14 × 20
3	1D	InSb	3.02–3.52	14 × 20
4	1E	InSb	3.52–4.08	14 × 20
5	2A	Si : Ga	4.08–5.30	14 × 20
6	2B	Si : Ga	5.30–7.00	14 × 20
7	2C	Si : Ga	7.00–12.40	14 × 20
9	3A	Si : As	12.40–16.50	14 × 27
10	3C	Si : As	16.50–19.50	14 × 27
11	3D	Si : As	19.50–27.30	14 × 27
12	3E	Si : As	27.30–27.70	20 × 27
13	4	Ge : Be	27.70–45.39	20 × 33

2.2. *Processing a Spectral Segment*

The *swsmake* code must correct for the deviations among the 24 spectra in each segment in order to combine them in a meaningful way. A number of effects prevent the off-line processing pipeline from completing this task. Several authors have described the dominant problem, known as the memory effect (e.g., Fouks 2001; Kester, Fouks, & Lahuis 2001). In the Si : Ga detectors of Band 2 (4.08–12.0 μm) and the Ge : Be detectors of Band 4 (27.7–45.4 μm), the dark current varies during a scan, depending on the recent flux history of the detector. OLP version 10.0 introduced an algorithm to model and correct for this effect, called *dynadark*, which is based on the Fouks-Schubert formalism (Fouks & Schubert 1995) and eliminates most of the problem in Band 2. Neither OLP 10.0 nor OLP 10.1 apply this correction to Band 4; even in Band 2, a sudden change in incident flux, such as might happen at the beginning of a scan of a bright source, results in a nonlinear response which *dynadark* cannot correct. Figure 2 illustrates what *dynadark* can accomplish in Band 2.

The interdependence of gain corrections and dark current corrections presents other problems. The relative spectral response function (RSRF) corrects all spectra for detector-to-detector variations in gain and variations of instrument throughput and detector response as a function of wavelength. If the additive correction for the dark current is poor, the multiplicative correction for the RSRF will generate spectral artifacts where the RSRF changes rapidly with wavelength. This problem is most obvious in Band 4, where *dynadark* is not applied, but artifacts also appear frequently in Band 2. The improved memory model in OLP 10 has reduced the artifacts in Band 2 significantly (see Fig. 2 in

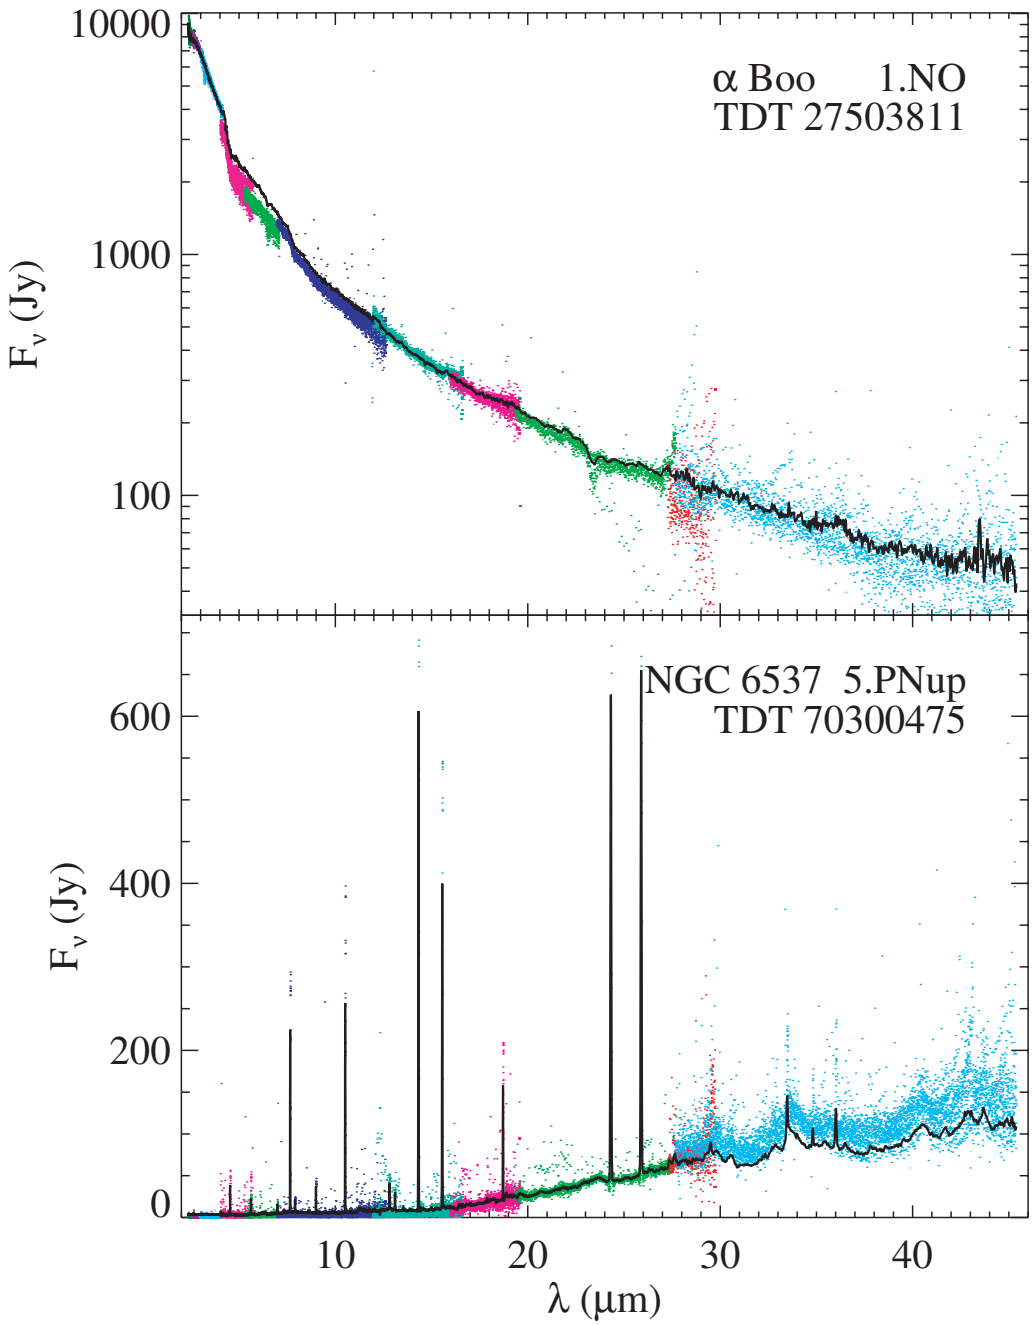

FIG. 1.—Sample AAR data (*colored dots*) for α Boo (*top*) and NGC 6537 (*bottom*), along with the results of our post–pipeline processing (*solid line*) for each spectrum. Data from different spectral segments (see text or Table 1) are shown in different colors. The infrared spectral classification (Paper I) and TDT number appear in the upper right corner of each plot.

Sloan et al. 2001). A similar problem occurs when the pipeline corrects for the detector-to-detector differences in gain. An additive error in dark current, which often varies with wavelength, can change the relative response of the detectors and distort their behavior with wavelength.

For these and other reasons, the spectra from individual detectors and scans in each spectral segment (even in Bands 1 and 3) must be aligned with one another in some way before they can be combined. The *swsmake* code performs this task by finding a filtered median spectrum for all 24 spectra (12 detectors, two scan directions), then correcting the individual spectra to fit the median. The objective of this step is to preserve the detailed spectral structure in each spectrum while removing the overlying deviations caused by variations in dark current or responsivity. As a consequence, the multiplicative correction is a simple cubic polynomial function of wavelength. For fainter spectra, the corrections are allowed to vary with wavelength only linearly, due to the poor signal/noise ratio in these data, and the correction is additive, since dark current will dominate the signal.

This shift from a multiplicative to an additive offset reflects the fact that at lower signals, potential errors in dark current subtraction will dominate multiplicative errors like

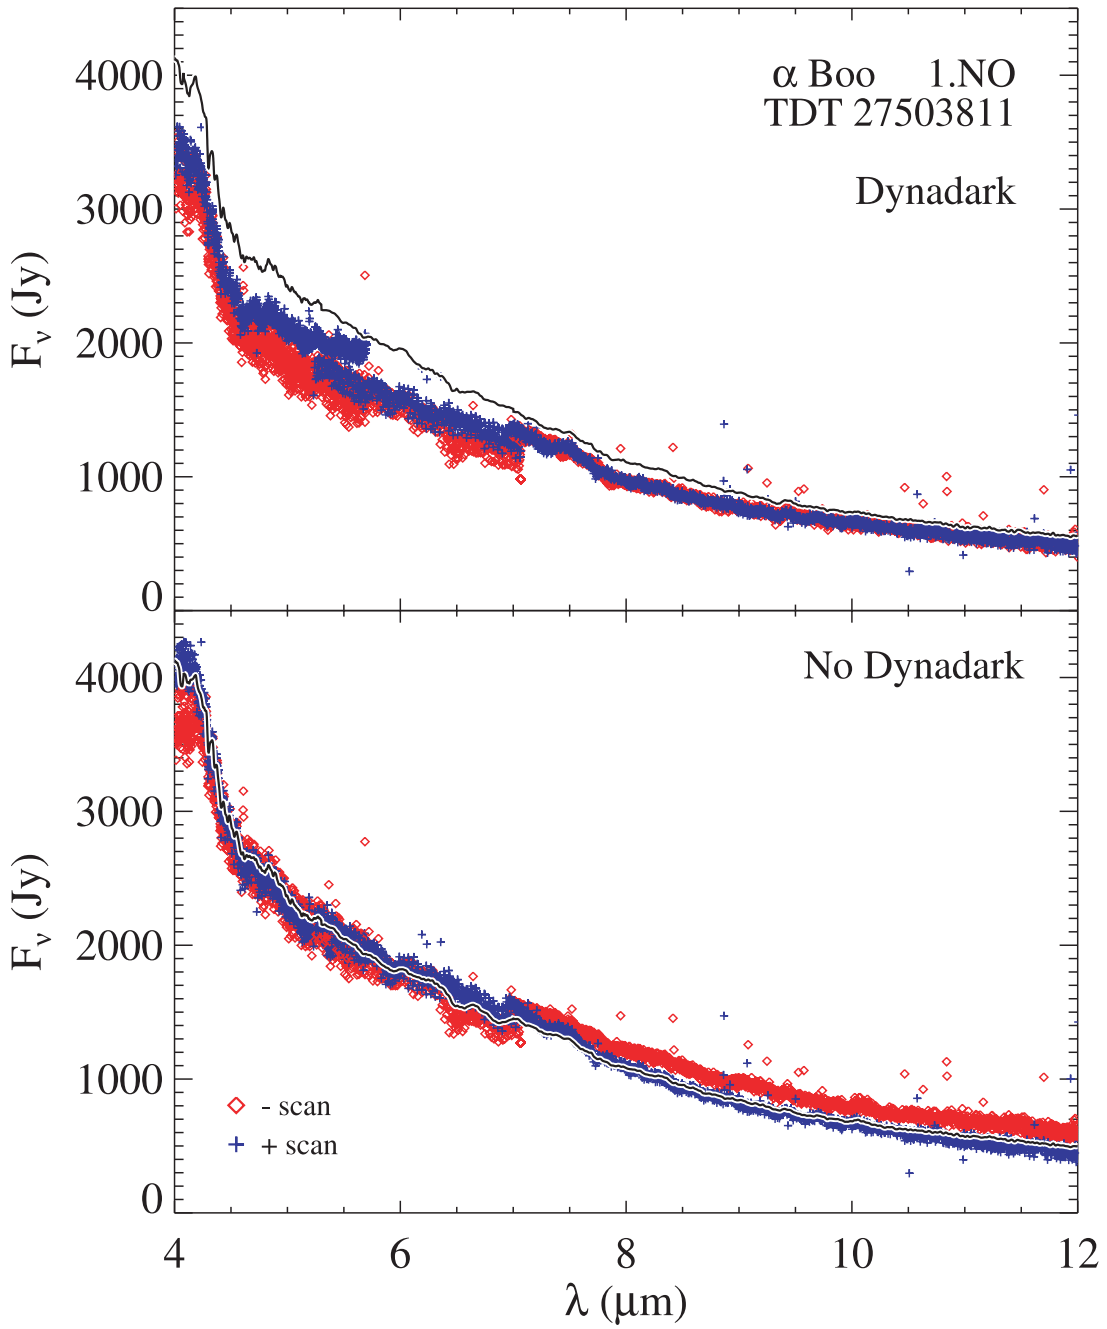

FIG. 2.—Spectra for α Boo from Band 2 (segments 5–7) processed with and without *dynadark* (*top and bottom*, respectively). The " $-$" scan direction appears as red diamonds, the " $+$ " direction as blue plus signs, and the final (normalized) spectrum as a solid line. The use of *dynadark* clearly improves the data in Band 2C (7–12 μm). In Band 2B, only the slope has changed slightly. In Band 2A (4.1–5.3 μm) *dynadark* has increased the mismatch between the scan directions, but it is actually attempting to improve the mean of these.

gain and responsivity differences among the detectors. By experimenting with spectra with different flux levels, we determined that the corrections should be additive if (1) all fluxes in the segment in question fell below 20 Jy, or (2) any fluxes fell below 15 Jy. The correction will be additive or multiplicative for an entire spectral segment, both when fitting spectra from individual detectors and scan directions to the median and when normalizing one segment to its neighbors (§ 2.3 below).

The spectra from individual detectors undersample the spectral resolution element of the SWS, but when combined, they interleave and oversample it. Thus, the spectrum from one detector might only catch the wings of an atomic line, while the spectrum from another detector might contain the core. To avoid undue influence of lines on the polynomial correction for individual detectors, the median is filtered so that it follows an approximate continuum. Figure 3 illustrates the generation of a median spectrum from the two scan directions for Band 1B of a typical star and its relation to the final spectrum for that segment. This method will not correct for a problem which affects all detectors simultaneously, such as telescope drift during a scan.

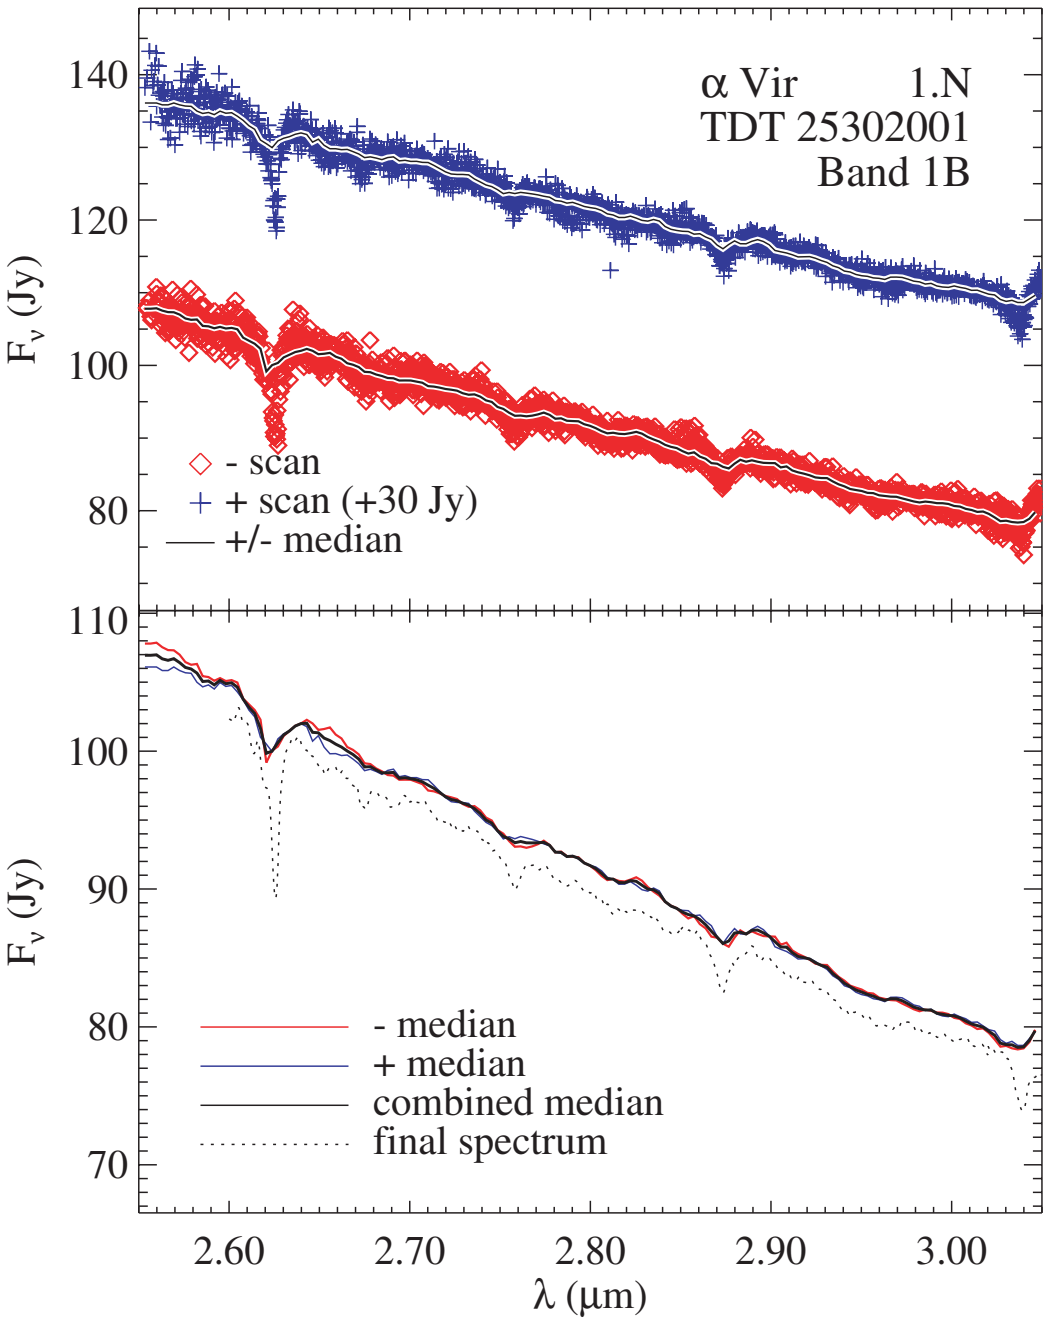

FIG. 3.—Band 1B spectrum for α Vir, showing the construction of a filtered median (*top*) and the final spectrum for the segment (*bottom*). In the upper panel, red diamonds depict raw down-scan data, blue plus signs depict raw up-scan data (offset vertically by 30 Jy). The solid black lines show the medians for each scan after removing atomic lines and other narrow features. The lower panel compares the medians from each direction (*in color*) and shows the combined median (*solid black line*) and the final (normalized) spectrum for the segment (*dotted black line*). Note that the correction of data from each detector to the median leaves the atomic lines intact.

The *swsmake* code combines all 24 corrected spectra, sorts by wavelength, flags bad data (using a sigma-clipping algorithm, discussed more fully in § B1), and regrids the surviving data onto a uniformly spaced grid using standard wavelength intervals based on the speed of the scan and the segment being processed (see Table 2). The new wavelength spacing is always on the order of the oversampled spacing of the original data. This regridding sacrifices some spectral resolution in order to gain uniformity among spectra in the database. The software generates formal uncertainties in the mean at each wavelength by comparing data at similar wavelengths from the original 24 uncorrected spectra.

2.3. *Combining Spectral Segments*

2.3.1. *Causes of Discontinuities*

The output from the pipeline often contains discontinuities between the spectral segments (e.g., Shipman et al. 2001; Sloan et al. 2001). The discontinuities arise from multiple causes. While the size of the discrepancies in Band 1 are

TABLE 2
SPACING IN ADOPTED WAVELENGTH GRIDS (IN NANOMETERS)

Segment	Band	Speeds 1 and 2	Speed 3	Speed 4
1.....	1A	0.50	0.25	0.125
2.....	1B	0.67	0.33	0.167
3.....	1D	0.67	0.33	0.167
4.....	1E	1.00	0.50	0.250
5.....	2A	1.00	0.50	0.250
6.....	2B	2.00	1.00	0.500
7.....	2C	2.00	1.00	0.500
9.....	3A	3.33	1.67	0.833
10.....	3C	3.33	1.67	0.833
11.....	3D	6.67	3.33	1.667
12.....	3E	6.67	3.33	1.667
13.....	4	10.00	5.00	2.500

typically a few percent, they can be substantial at longer wavelengths.

The discontinuities generally arise because the apertures truncate some flux from the source. The aperture size increases between Bands 2C and 3A (12.0 μm), Bands 3D and 3E (27.5 μm), and Bands 3E and 4 (27.7 μm) (see Table 1). For an extended source, such as an H II region or planetary nebula, the discontinuity can be quite dramatic.

Discontinuities can also appear in spectra of point sources, because the size of the point spread function (PSF) is similar to the size of the aperture. Figure 4 plots the relative size of the aperture dimensions, the full width at half-maximum (FWHM) of the PSF, and diameters containing 92%–93% of the light from a point source as a function of wavelength. The PSF is much smaller than the aperture in Band 1, and corrections for this band tend to be small. As the wavelength increases, though, a larger fraction of the flux falls outside the aperture even when the source is well-centered. Any pointing error, whether it arises from the absolute pointing of the telescope, an error in the coordinates given by an observer, or drift of the telescope during a spectral scan, will decrease the amount of flux incident on the detector.

The pointing model improved through the *ISO* mission so that the final error was typically less than $\sim 1\farcs4$ (2 σ; Salama et al. 2001). Errors in input coordinates, though, can be as large as several arcseconds, as discussed in Paper I. Internal *ISO* reports show that pointing errors of $5''$ reduce the flux levels by $\sim$20% in Band 1 and $\sim$10% in Band 2 (Salama 1998). The PSF is significantly larger compared to the aperture size at longer wavelengths. For example, the diameter containing 92% of the flux at 40 μm is $26''$, compared with the $19''$ aperture (short axis). Thus, a substantial fraction of the flux is already truncated, and a $5''$ shift produces a smaller error ($\sim$5%).

Drift in the telescope pointing during a scan can produce discontinuities at boundaries between bands where the size of the aperture does not change. Typical drift is smaller than $\sim 0\farcs7$ (Salama et al. 2001), but this is enough to add $\sim$2%–3% offsets (either random or systematic) to the data (Salama 1998). For a well-centered source, the effect of drift is small since the source must shift a significant distance before a noticeable amount of flux is truncated by the aperture, but if the source is already off center (due to absolute pointing error or poor coordinates), then the effects of drift are enhanced. These pointing errors are a significant cause of the discontinuities between segments (Shipman et al. 2001).

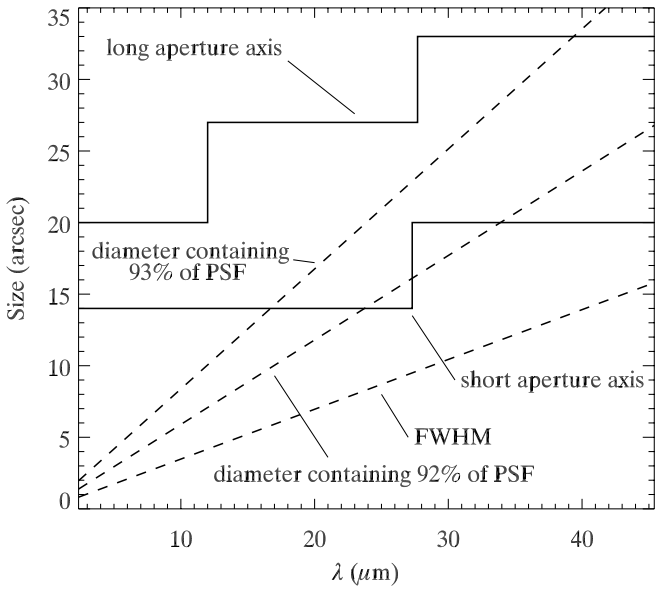

FIG. 4.—Size of the SWS apertures as a function of wavelength for the long axis (*upper solid curve*) and short axis (*lower solid curve*). The dashed lines show the FWHM and the diameters containing 92% and 93% of the flux from a diffraction-limited PSF from a 60 cm telescope with a 20 cm central obscuration, giving an indication of the relative sizes of the PSF and aperture. The FWHM for this PSF contains 73% of the flux.

Because the size of the PSF changes continuously with wavelength, it varies within a spectral segment as well as from one segment to the next. Drift or random jitter in the telescope pointing during a scan can have a similar result. This behavior produces discontinuities in both the total flux and the slope of the spectrum, and it can distort the shape of the spectrum within a segment. As discussed below, our reprocessing algorithm corrects for the discontinuities in total flux, but without detailed knowledge of the precise location of the source within the aperture during a scan, the algorithm cannot correct for discontinuities in slope between segments or distortions within a segment.

2.3.2. *Correcting for Discontinuities*

The *swsmake* software combines the 12 spectral segments into a single spectrum by normalizing segments to their neighbors to eliminate discontinuities and eliminate overlap between segments. Figures 5 and 6 compare the starting point and endpoint of this process for a spectrum of α Vir.

Either Band 1E (segment 4, 3.52–4.08 μm) or Band 3D (segment 11, 19.5–27.3 μm) serves as the starting point, depending on which contains more flux. Normalization begins at one of these bands and propagates outward, band by band. Any errors in the absolute flux level within the band chosen as a starting point will be propagated to the other bands through the normalization process.

For each pair of neighbors from Band 1A to 3D, the algorithm determines a correction by comparing the average flux in the region of overlap (as defined by Table 3). For bright sources the normalization factor is multiplicative and is the ratio of the averages. For faint sources, the normalization is additive and found from the difference in averages.

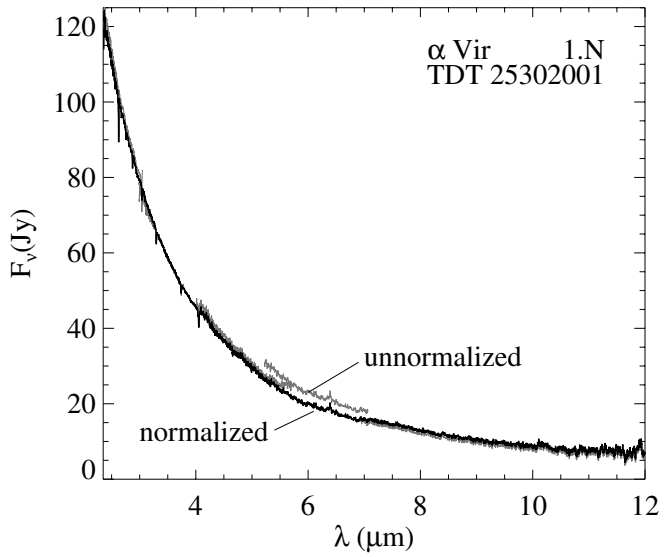

FIG. 5.—Bands 1 and 2 for α Vir comparing the unnormalized data (*gray lines*) to the data after final normalization (*black line*).

For any pair of neighboring segments, the SWS documentation (e.g., Leech et al. 2002) considers data in the overlap region to be " in-band " for only one of the two segments, but we exploit the fact that some of the data in these overlap regions is valid for both segments. We have chosen the overlap regions (Table 3) to avoid poorly behaved data in either spectral segment. While it is true that the band edges occur in spectral regions of low detector responsivity, in general the data in both bands are suitable for normalization within the wavelength ranges we have specified. The notable exception is the interface between Bands 2C and 3A, where artifacts from the RSRF often appear at 12.3 μm in Band 3A (see § 5.5) and difficulties with the memory effect in bright sources can lead to inaccurate RSRF corrections in Band 2C (see § 5.2).

The normalization from Band 3D to Band 3E to Band 4 (segments 11–13) requires more care. The long-wavelength end of Band 3D is often contaminated by a light leak which can turn the spectrum up or down near 27 μm. The band is rarely trustworthy beyond 27.3 μm as the detector response is quite poor at these wavelengths. Band 3E typically contains poor data and even when good, the data are questionable longward of 27.7 μm. We have found that the data in Band 4 are reasonable down to 27.5 μm, but this still does not close the gap to Band 3D. Since there is no overlap, the software must extrapolate over the gap to determine the normalization factor. Our procedure is to extrapolate to fit Band 4 to Band 3D and then fit Band 3E to the now-normalized Bands 3D and 4. Section B2 provides the details.

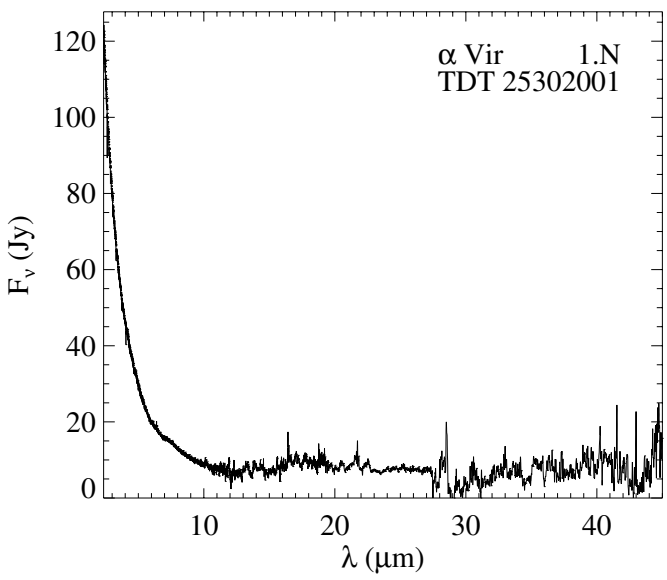

FIG. 6.—Full set of data for α Vir after final normalization

TABLE 3
NORMALIZATION REGIONS

Segments	Bands	λ Range (μm)
1–2	1A–1B	2.562–2.630
2–3	1B–1D	3.000–3.043
3–4	1D–1E	3.351–3.670
4–5	1E–2A	4.019–4.082
5–6	2A–2B	5.265–5.400
6–7	2B–2C	7.02–7.06
7–9	2C–3A	12.25–12.55
9–10	3A–3C	16.00–16.56
10–11	3C–3D	19.37–19.56
11–12	3D–3E	No overlap
12–13	3E–4	27.70–28.80

The extrapolation process is sensitive to the detailed shape of the spectrum in the 20–35 μm region, and this has forced us to depart from applying a standard algorithm to all sources. Instead, the extrapolation method depends on the shape of the spectrum, as determined by the classifications in Paper I and explained in § B2.

Strong emission lines in some spectra can affect the normalization if they occur in the overlap regions used to normalize segments. As described in § B2, these are also treated as a special case to the standard normalization procedure.

When normalizing two adjacent segments, the *swsmake* code estimates the uncertainty in the correction and propagates this into the uncertainty analysis. The spectral data files generated by the *swsmake* code contain two columns of uncertainties. The first column is the statistical uncertainty in the mean found when combining the data from all detectors and scans in a given segment into one spectrum (see § 2.2), essentially the point-to-point or spectroscopic uncertainty. It depends only on the range of data at a given wavelength and is not a measure of any systematic errors in the data. The second uncertainty estimate at each wavelength includes the uncertainties arising from our normalization procedure combined with the uncertainty in the mean. We will refer to it as a segment-to-segment or normalization error. Again, these estimates are based on the range of the data in the overlapping regions and do not estimate systematic effects (which, for example, could result from the interplay of dark current and gain on our additive or multiplicative corrections). These normalization uncertainties are propagated from one segment to the next, with the exception of Bands 3E and 4, where the extrapolation prevents a clear estimate; for these bands, the uncertainty in Band 3D is used. This uncertainty is not a formal

photometric error because it makes no attempt to estimate the absolute amount of flux truncated by the aperture edges; it only attempts to estimate the relative segment-to-segment uncertainty in the normalization process.

The final step in the *swsmake* algorithm is to eliminate the overlap between segments by adopting the wavelength ranges in Table 1 for each segment. These ranges differ slightly from the in-band ranges as defined by Leech et al. (2002), particularly for Bands 3D, 3E, and 4, and at the boundary between Bands 2C and 3A. Data outside the range for a given spectral segment are truncated. The result is a smooth, continuous spectrum from 2.4 to 45 μm, as shown in Figure 6.

3. THE ATLAS

Reprocessed spectra of all valid SWS full-scan observations are available as an atlas electronically as explained in Appendix A. For each spectrum, there are three files: the results of the reprocessing of individual spectral segments (suffix "pws"), the results of normalizing the segments into a continuous spectrum (suffix "sws"), and the log file (suffix "log"). All three files are simple ASCII text files.

Both types of data file ("pws" and "sws") have the same format. Each begins with two lines of 12 numbers. The first line gives the total number of wavelength elements for each spectral segment. The second gives a number for each segment that indicates whether the corrections and normalization for that segment are multiplicative (0) or additive (1). The data follow these two header lines, one line for each wavelength element. The first column is wavelength in μm, the second is flux density in Jy, the third is spectroscopic uncertainty in the flux density, and the fourth is normalization uncertainty (both in Jy). Since the latter uncertainties are not calculated until the segments are normalized to each other, the fourth column is equivalent to the third in the "pws" files. These identical formats allow both spectral data files to be read with the same software. The log files provide details about the decisions made by the software.

The spectral atlas organizes files by the TDT number that identifies the *ISO* observation. A complete catalog of observations, with the TDT number and organized by coordinates, appears in Paper I. This catalog also contains classifications, allowing a search for all spectra with similar spectral properties.

The spectra observed by the SWS in AOT 1 and AOT 6 total 1271 (1248 and 23, respectively). Of these, nine AOT 1 spectra are sufficiently flawed to prevent the pipeline from producing an AAR file, even though a browse product was available at one time (browse products are no longer available for these flawed spectra). Paper I classified six as Group 7 (fatally flawed); the remainder are TDT 12601620 (an offset pointing near γ Dra classified as 1.NO), TDT 14801733 (AFGL 2094, 5.UE), and TDT 40201614 (an offset spectrum near Sharpless 184 classified as Group 6). As a result, the atlas contains 1262 spectra.

4. VALIDATION

There are a number of ways to assess the validity of the post–pipeline processing by the *swsmake* code. Below, we compare the *swsmake* results to (1) synthetic spectra and previously calibrated spectra of standard stars, (2) published SWS spectra processed by other scientists, (3) SWS AOT 6 spectra processed with a variation of the *swsmake* code, and (4) browse products from IDA.

4.1. *Standard Stars*

Figure 7 compares the processed spectra of two infrared standard stars, α CMa (A1 V) and γ Dra (K5 III), with the Kurucz model for α CMa presented by Cohen et al. (1992) and the absolutely calibrated composite spectrum of γ Dra by Cohen et al. (1996). In both cases the SWS spectrum produced by *swsmake* agrees well with the reference spectrum. Most of the discrepancies lie at longer wavelengths where the low fluxes limit the quality of the data.

Figure 8 compares the same spectra, but with the vertical axis plotted in $\lambda^2 F_\nu$ units (equivalent to $\lambda^4 F_\lambda$). This exaggerates problems at longer wavelengths, but it allows a more detailed comparison of the continuum shape and the shape and strength of the atomic and molecular features. The SWS spectrum of γ Dra agrees quite well with the composite spectrum of Cohen et al. (1996), although there are some detailed differences to the shape of the combined SiO overtone and CO fundamental at 4–6 μm and the SiO fundamental at 8–10 μm.

The deviations between the spectrum of α CMa and the calibrated model of Cohen et al. (1992) are more significant. The dominant difference is the upswing in the SWS data at 12 μm, a result of a less-than-successful correction of the dark current memory problem by *dynadark*. It is noteworthy that the artifacts seen in early-type spectra at 4.5 and 8 μm by Price et al. (2002) produced by OLP 10.0 are much weaker or absent in the OLP 10.1 data. The overall

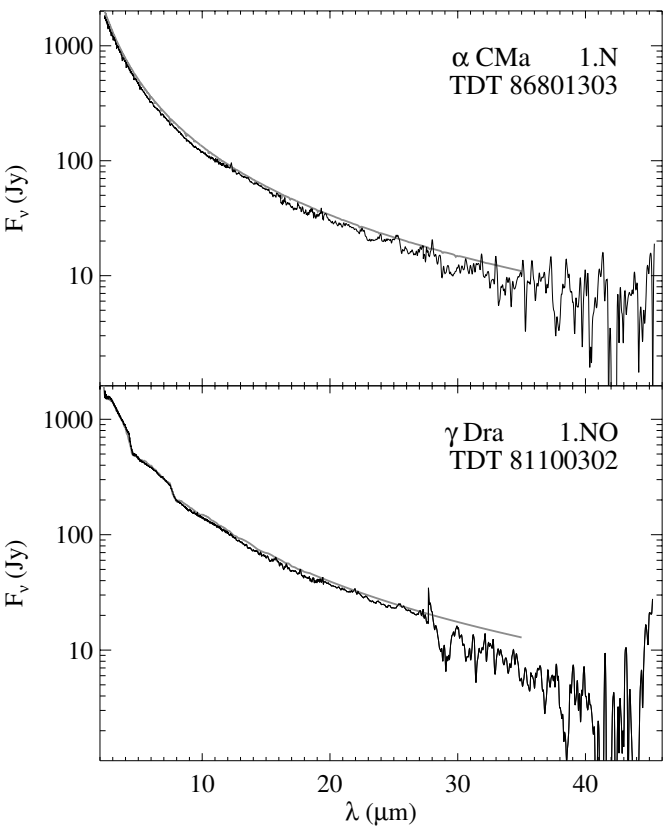

FIG. 7.—Comparison of the *swsmake* results (*black lines*) for the standard stars α CMa (*top*) and γ Dra with data from Cohen et al. (1992) and Cohen et al. (1996) (*gray lines*).

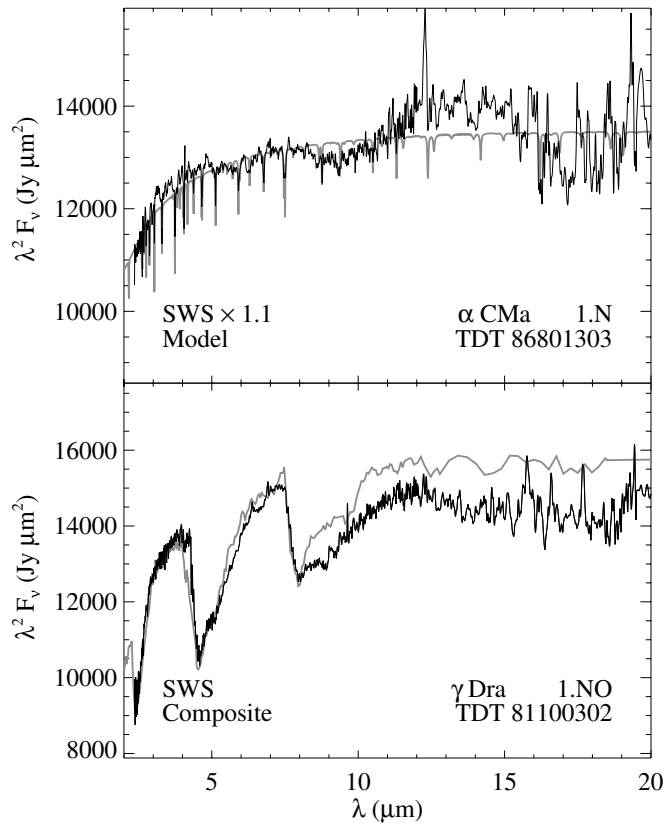

FIG. 8.—Same as Fig. 7, except that the vertical axis is plotted in $\lambda^2 F_\nu$ units to better show deviations from a featureless stellar continuum.

flux level of our post-pipeline SWS spectrum is lower compared to the model data by a factor of 1.1 (in the 2.4–10.0 μm range), and given the nonphotometric nature of the SWS, this should not be surprising. (The other full-scan spectrum of α CMa obtained by the SWS, TDT 68901202, is brighter than the calibrated model by roughly the same amount.) Most other standards fare similarly in such a comparison.

4.2. *Comparison with Other Processing Methods*

Most published SWS spectra use earlier versions of the pipeline, making it difficult to distinguish between differences due to changes in the pipeline and those caused by differences between interactive processing and the automated *swsmake* algorithm. Vandenbussche et al. (2002) and Onaka et al. (2002) presented several spectra of stellar sources processed from OLP 10 pipeline results. To increase the sample of spectra for comparison, we also examine results from earlier versions of the pipeline from other authors.

Vandenbussche et al. (2002) included a number of speed 4 spectra taken during the main part of the *ISO* mission in an atlas primarily devoted to Band 1 spectra (2.35–4.05 μm) from the postcryogenic phase of the mission. Figure 9 compares their results for five speed 4 spectra with the corresponding *swsmake* results. Their spectra started with OLP 10.0 results, but for Band 1 this is identical to OLP 10.1. Overall, the two sets of data agree closely. Table 4 lists the scaling factors used to normalize the segments by *swsmake* and Vandenbussche et al. (2002). The average difference in the scaling factors is only 0.002, or less than 1%. The minor differences in the spectra result from differences in smoothing and definitions of out-of-band data. The emission lines are stronger in δ Cen and the absorption lines deeper in the other stars in the spectra from Vandenbussche et al. (2002) because the *swsmake* code smooths the data slightly when it regrids them. On the other hand, our spectra are considerably less noisy from 3.0 to 3.7 μm, probably because we discard most of the out-of-band data from Bands 1D and 1E and they do not.

TABLE 4
COMPARISON OF NORMALIZATION CONSTANTS

Star	Method	1A	1B	1D	1E
α Vir	*swsmake*	0.973	0.979	1.001	1.000
	Vdb	1.00	1.00	1.02	1.00
α CMa	*swsmake*	1.021	1.016	1.014	1.000
	Vdb	1.00	1.00	1.00	1.00
γ Dra	*swsmake*	0.996	1.002	1.002	1.000
	Vdb	1.00	1.00	1.00	1.00
γ Cru.....	*swsmake*	1.201	1.246	0.999	1.000
	Vdb	1.20	1.22	1.00	1.00
δ Cen.....	*swsmake*	0.983	0.974	0.995	1.000
	Vdb	1.00	1.00	1.00	1.00

NOTES.—"Vdb" refers to Vandenbussche et al. 2002. The exact correspondence in Band 1E reflects the fact that this band is held fixed in our normalization of stellar sources and the fact that Vandenbussche et al. 2002 attempts to hold all bands fixed if possible.

Onaka et al. (2002) observed Z Cyg, an oxygen-rich M5 star, as part of a study of infrared variability in circumstellar dust emission. Figure 10 compares our spectra with the same observations reduced by Onaka et al. (2002) from OLP 10.1 pipeline output. As with the Band 1 data of Vandenbussche et al. (2002), the two reduction methods produce generally indistinguishable results. The two sets of data resemble each other closely even in the problematic Bands 3E and 4, which suggests no serious flaws with either processing method. The discrepancies at the longer wavelengths for TDT 26300316 occur because the normalization for the brighter bands was multiplicative, but the fainter fluxes in Bands 3E and 4 triggered a switch to additive offsets. This switch does not automatically result in discrepant results, however, as can be seen in the other Z Cyg spectra where the same switch did not affect the resulting spectral shape.

Sturm et al. (2000) examined the emission from four galaxies: M82, NGC 253, NGC 1068, Circinus, and the 30 Dor region of the Large Magellanic Cloud. Although they started with an earlier version of the pipeline, a comparison with the *swsmake* results is still instructive (Fig. 11). Generally, the spectra agree well in Band 3 (12–27 μm). At longer wavelengths, the differences arise from the fact that they do not normalize Band 4 to Band 3. The differences at the shorter wavelengths also result from the normalization process. The galaxies, as with most spectra in groups 4 and 5, are relatively faint below 10 μm compared to the longer wavelengths. As with Z Cyg, this change in flux level triggers a change in the normalization from multiplicative to additive. While this change did not affect 30 Dor significantly, its effects can be seen in the spectra of the four galaxies. Sturm

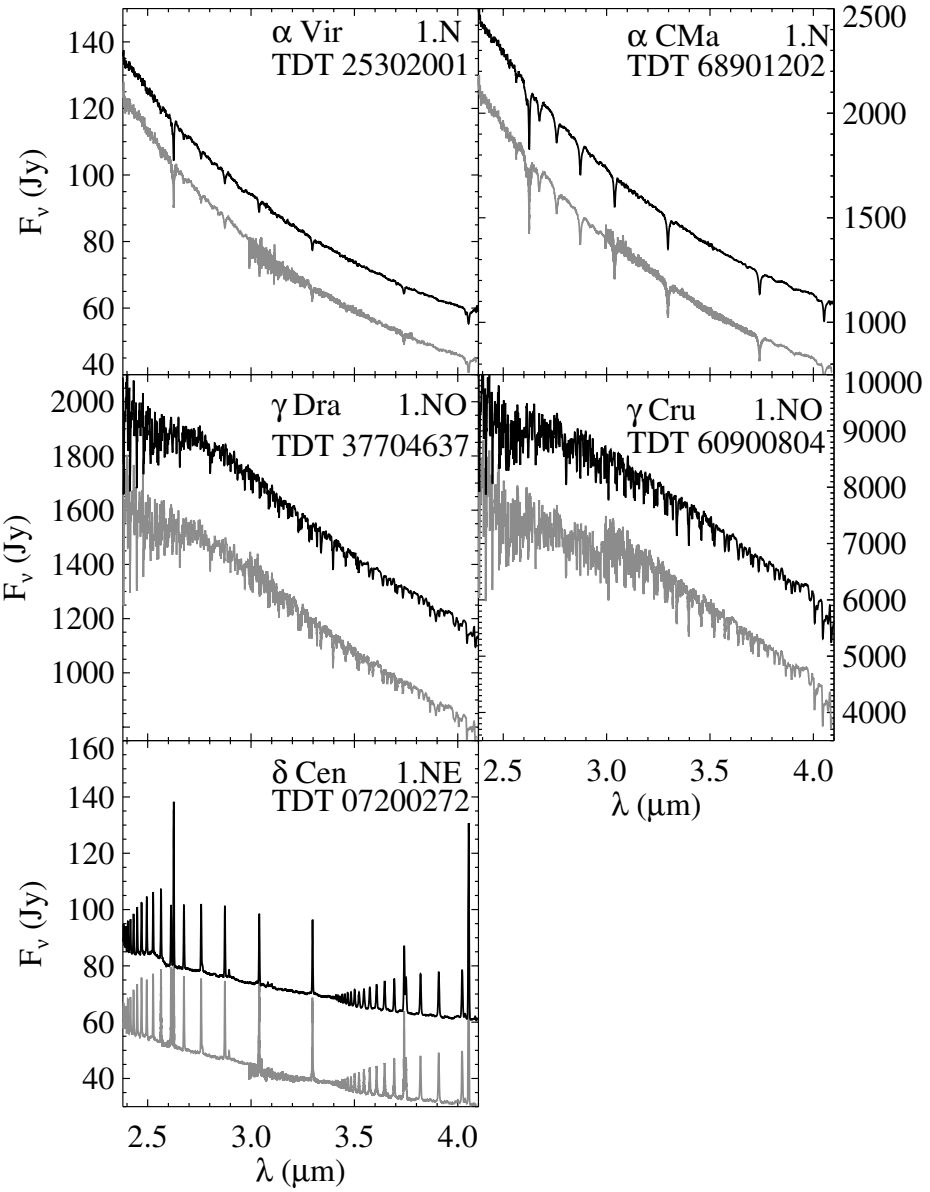

FIG. 9.—Comparison of results from *swsmake* with those of Vandenbussche et al. (2002) in Band 1. The *swsmake* spectra (*black lines*) have been offset from the comparison spectra (*gray lines*) for ease of comparison: α Vir +15 Jy; α CMa +300 Jy; γ Dra +350 Jy; γ Cru +1500 Jy; δ Cen +30 Jy.

et al. (2000) processed their data before *dynadark* became a part of the pipeline, so uncorrected memory effects could also influence the shape of Band 2 in their data. At these low flux levels, though, memory effects are usually not significant.

In general, the fainter the flux in a spectrum, the less trustworthy the normalization becomes. This happens at short wavelengths in spectra of galaxies and H II regions, and it also happens at the red end of stellar spectra. In these cases, it may be better to use the unnormalized data at certain wavelengths as found in the "pws" file for a given source.

4.3. *Comparison with AOT 06 Data*

Twelve of the sources observed by the SWS using AOT 1 were also observed using AOT 6, an observing mode designed for higher spectral resolution and sensitivity. Although the software used to reduce the AOT 6 observations was adapted from the *swsmake* code (see App. C), the differences in the observing modes make comparison of the common sources useful. While each pair of spectra comes from the same astronomical source, differences could arise from the different observing modes or variations in the instrumental response. For example, the memory effect depends on background and dwell time, and the AOT 1 and 6 scans probe different regimes in both parameters. Figure 12 shows the unnormalized and normalized spectra for three representative sources, W51 IRS 2, NGC 7027, and V CrB. The similarities in the spectra demonstrate the robustness of the *swsmake* data processing algorithm, as well as that of the SWS instrument itself. Comparisons of other spectra give similar results.

The normalized spectra match each other well, particularly Bands 1, 2, and 3 for W51 IRS 2 and NGC 7027. The differences in Band 4 arise primarily from the uncorrected memory effect in the Ge : Be detectors (see §§ 2.3 and 5.1),

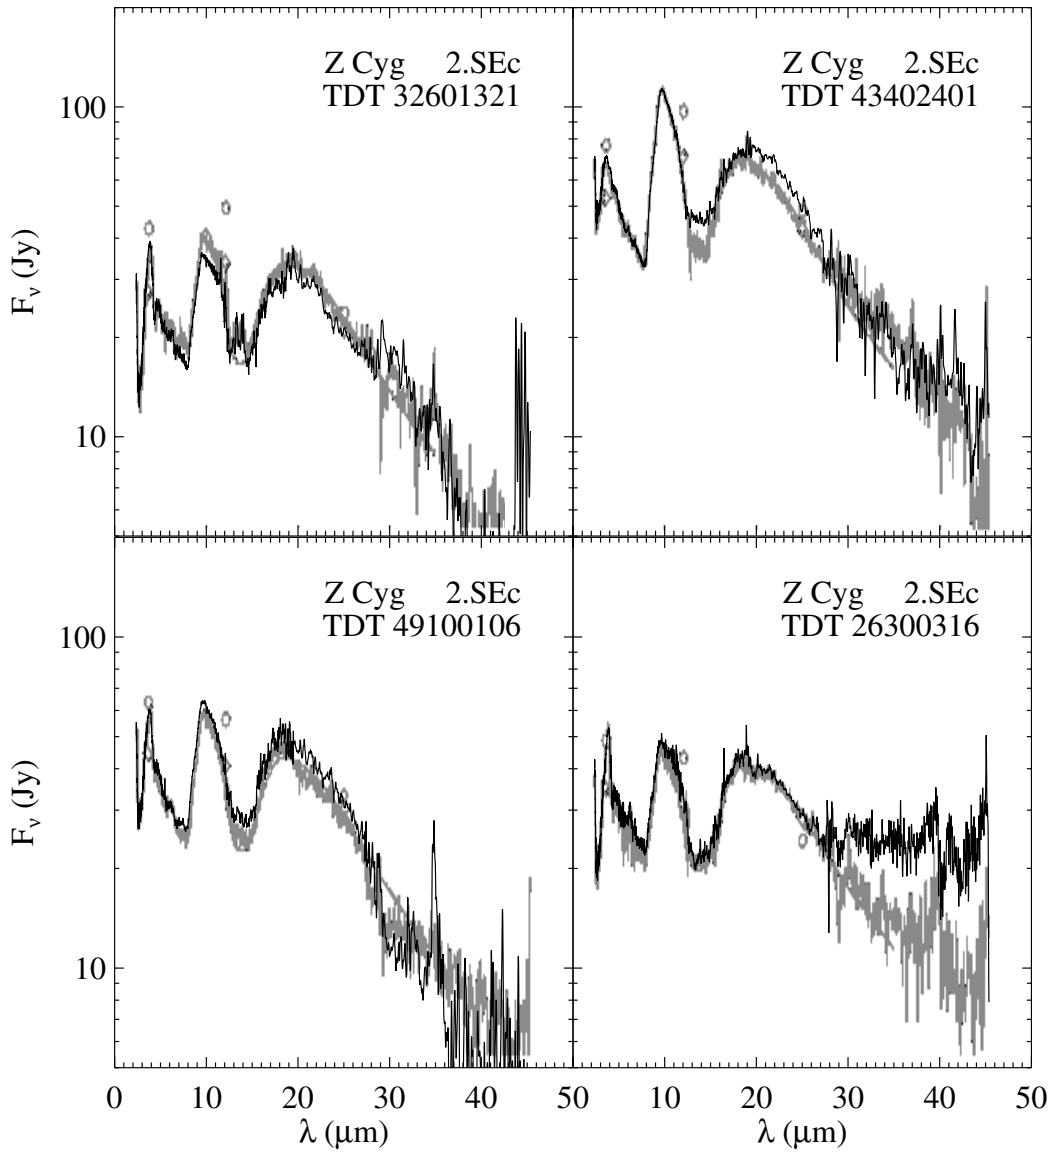

FIG. 10.—Comparison of results from *swsmake* (*black lines*) with those of Onaka et al. (2002), adapted from their Figure 1 (*gray lines*) for the evolved variable star Z Cyg. The two sets of data are virtually indistinguishable at most wavelengths. The small circles are their photometric measurements of Z Cyg.

and the often problematic interface between Bands 3D, 3E, and 4. At these bright flux levels, Band 3E in AOT 1 behaves poorly and can be difficult to salvage. The AOT 6 spectra often provide a better reference in this region, and because of the difficulty of normalizing AOT 1 data through this region, also present more reliable flux levels for most of Band 4.

All stellar sources observed in AOT 6 are variable stars, complicating the comparisons. Figure 12 compares AOT 6 and AOT 1 spectra of V CrB taken slightly more than one period apart. V CrB is a Mira variable with a period of 358 days, and the two observations are 421 days apart. This difference is 63 days more than one period, and the spectra are reasonably similar. The differences in the V CrB spectra are generally consistent with the variability of the source.

4.4. *Browse Products*

The browse products are automatically produced by scaling all reliable (i.e., unflagged) data from the individual detectors to each other, eliminating outliers using sigma clipping, and rebinning the result to a more uniform grid. In principle, these steps produce results analogous to our "pws" files (i.e., before we normalize the segments), but as Figure 13 illustrates, the two methods produce significantly different results.

Our postprocessing pipeline produces spectra with fewer spikes. Figure 13 shows MWC 1080 as an example of our detection of glitches that escaped the OLP deglitcher. This results primarily from the different methods of identifying spikes. The OLP pipeline identifies spikes from discontinuities in the signal ramps of individual detectors during an integration at a given wavelength. The *swsmake* software identifies spikes by comparing data from different detectors and scan directions. Thus, the software eliminates or significantly reduces spikes that the OLP pipeline failed to catch.

Browse product spectra often contain discontinuities between spectral segments. In the worst cases, these can have a dramatic effect on the appearance of the spectra. The

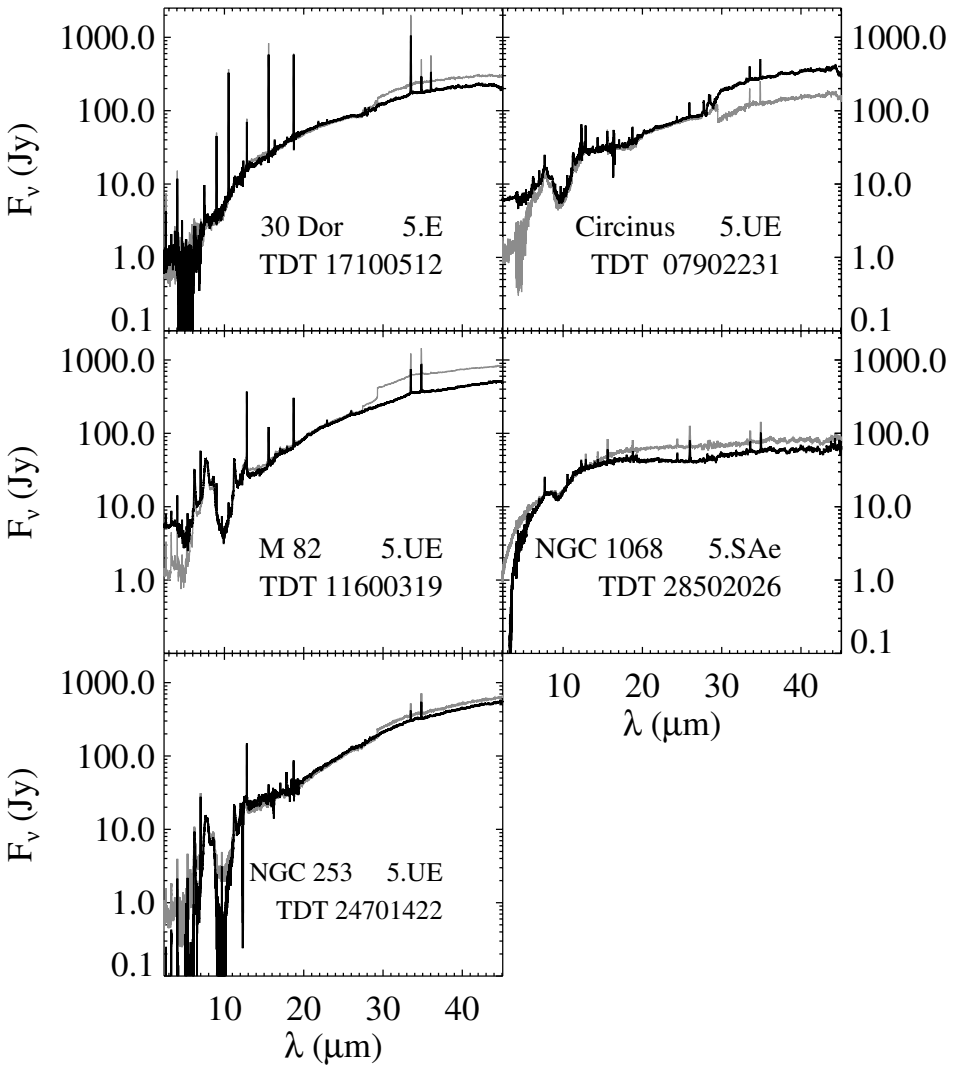

FIG. 11.—Comparison of results from *swsmake* (*black lines*) with those of Sturm et al. (2000) (*gray lines*) for five extragalactic sources

examples in Figure 13 show how our normalization procedure can reveal the true nature of a source which had been disguised by the discontinuities in the browse product. Indeed, many of the classifications made in Paper I proved impossible with just the browse product; preliminary versions of our post–pipeline processing were required to produce a more trustworthy overall shape. The elimination of the overlaps between segments in the browse products prevents the application of any simple segment-to-segment normalization to these data as done by the *swsmake* algorithm.

We should also point out that the agreement between *IRAS* photometry and our final SWS spectra is not very uniform, since many sources are variable and our choice of a spectral segment as the fiduciary is somewhat arbitrary. As an example, in Figure 13 our GY Aql spectrum is closer to the *IRAS* photometry than the browse product, but the opposite is true for AFGL 2199, and both versions agree for MWC 1080.

5. CAVEATS

While we have optimized the *swsmake* code to produce reliable results, the potential user should be aware that limitations still exist.

5.1. *Memory Effects*

The memory problem outlined in § 2.2 affects Bands 2 and 4. Since the implementation of *dynadark* in OLP 10.0, most Band 2 spectra are well behaved, but *dynadark* has not been applied to Band 4, so the overall shape of the spectra from 29 to 45 μm must be treated with caution. Furthermore, *dynadark* cannot always completely correct for hysteresis in the dark current for bright sources in Band 2. The impact is most pronounced at the long-wavelength edge of Band 2C (12 μm), where the poor responsivity and a poor dark current correction can cause the spectrum to turn up or down (see Fig. 14). Therefore, any kink in the spectrum at the boundary between Bands 2C and 3A ($\sim$12.4 μm) should be viewed with suspicion. Similarly, unexpected structure in Band 4 should be interpreted cautiously.

Changes in how data were obtained and calibrated during the mission had a strong impact on the memory effect, as discussed in § 5.7 below.

5.2. *Normalization Issues*

The normalization of spectral segments remains problematic for a number of reasons. Consequently, we include the

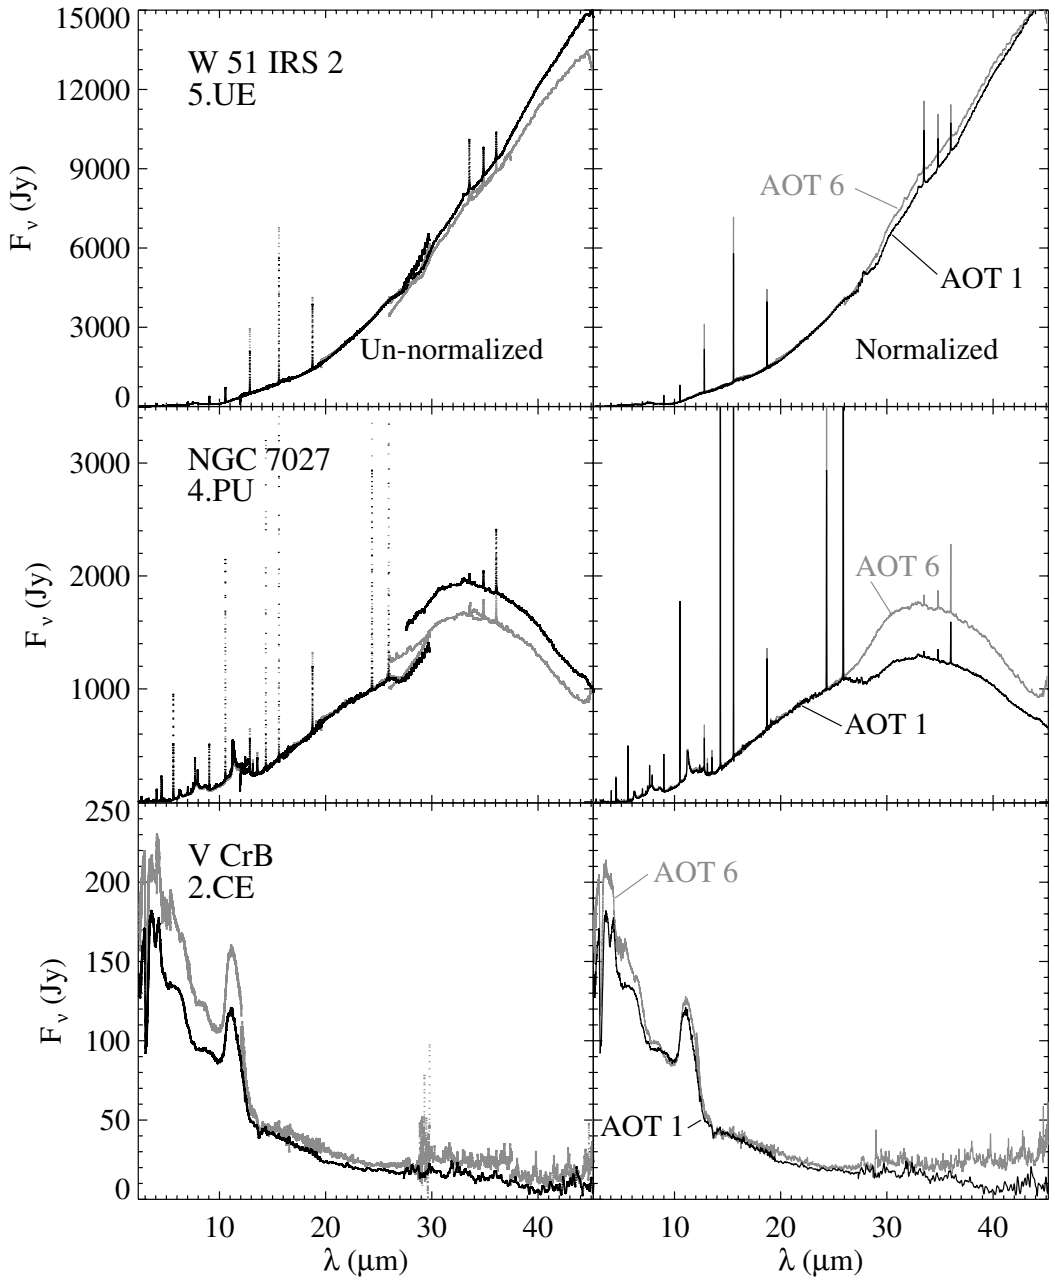

FIG. 12.—Comparison of the *swsmake* results for AOT 1 spectra (*black lines*) and AOT 6 spectra (*gray lines*) for three sources: W51 IRS 2 (*top*), NGC 7027 (*middle*), and V CrB (*bottom*). The TDTs used are W51 IRS2, 12801416 vs. 31901604 (AOT1 and AOT6, respectively), NGC 7027, 55800537 vs. 33800505, V CrB, and 67600104 vs. 25502252. V CrB is a Mira variable with a period of 358 days; the observations are 421 days, or 63 days more than one period, apart.

spectra both before and after segment-to-segment normalization in the atlas, allowing the user to choose which spectrum to adopt, or even to apply their own normalization algorithm.

For extended sources that overfill the aperture, the normalization over the boundaries where the aperture size changes is somewhat artificial. The spectra on either side of one of these boundaries actually arise from two physically distinct objects, since the spectrum from the larger aperture includes regions of the source not included in the spectrum from the smaller aperture. To be consistent, we applied the normalization algorithm to all the spectra.

In principle, models of how instrumental throughput onto the detectors varies with wavelength and source position within the aperture could be used to correct for pointing-induced throughput errors. However, the absolute positional uncertainty is at best on the order of $\sim 1\farcs4$, and this is enough to prevent an accurate correction. For this reason, we opted to make empirical segment-to-segment normalizations based on the data themselves. If truncation of the PSF or some other problem affects the starting point for the normalization process, then the errors introduced will propagate through the entire spectrum. Some spectra show reduced flux levels due to poor coordinates and other pointing-related issues, and the user can renormalize the entire spectrum if needed.

A more subtle problem involves variations in the throughput or gain within a segment, either as a result of the

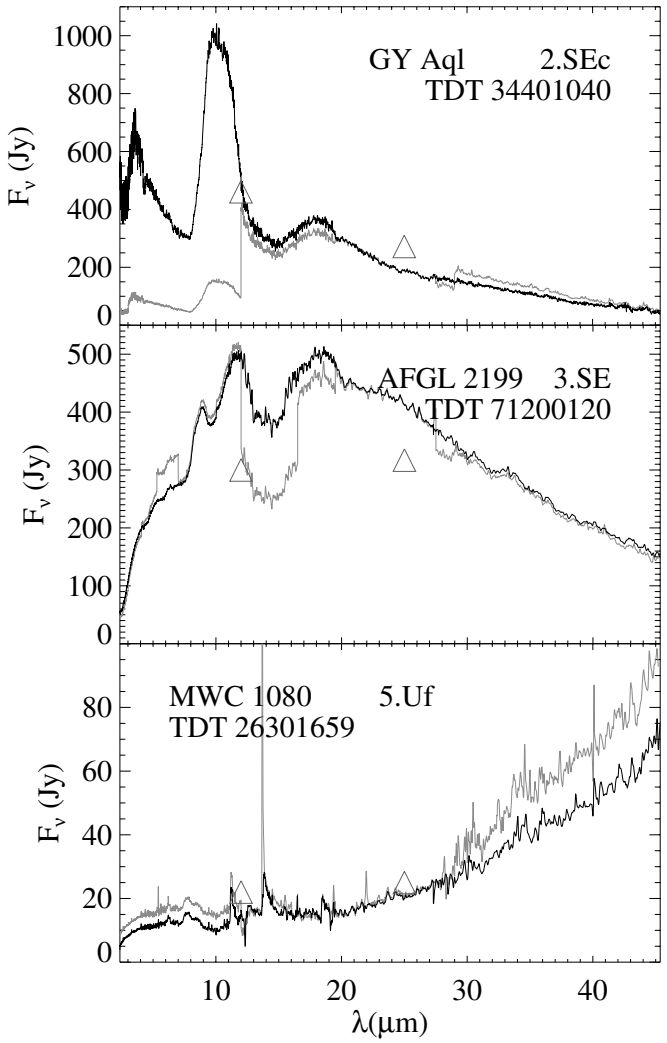

FIG. 13.—Comparison of the *swsmake* processing (*black line*) with the *ISO* browse product (*gray line*) and *IRAS* photometry (*triangles*) for three representative spectra. The top and center panels illustrate how our normalization process removes discontinuities which can disguise the true nature of a spectrum for a blue and red source, respectively. The bottom panel shows how our spike rejection process removes most of a strong spurious feature at 13.5 μm which appears in the browse product. Our algorithm also rejects several other spikes flagged by *ISO* which appear in the browse data (e.g., at 5.4, 19.3, and 22 μm).

dependence of the size of the PSF with wavelength or due to drift or pointing within a scan. These variations will modify the slope of the spectrum, and the *swsmake* algorithm does not address them.

Disentangling additive and multiplicative normalization corrections can also lead to problems. The segment-to-segment normalization assumes that the correction for an entire segment is either additive or multiplicative, when in reality it might possibly be both. As a result, the low-flux tails of some spectra, after normalization, might actually be negative, particularly at the long-wavelength end of stellar or near-stellar spectra (Groups 1 and 2), or at the short-wavelength end of very red spectra (Groups 4 and 5). Figure 11 illustrates this problem for the latter group, showing cases where the normalization overestimates and underestimates the short-wavelength flux. Approximately half of the spectra in Group 4 have some negative flux values in Band 1.

Finally, we must emphasize that the responsivity of the detectors is low in the regions of overlap of adjacent detectors. In general, no obvious problems attributable to this appeared in our analysis. However, when the memory effect and low responsivity at 12 μm in Band 2C have caused a turn-up or turn-down in the spectrum, this error will affect the normalization between Bands 2C and 3A (segments 7 and 9), and it will propagate to other segments in the normalization sequence (as Fig. 14 demonstrates).

5.3. *Emission Lines*

The OLP 10.1 pipeline successfully identifies most spikes in the data, and it largely corrects problems in previous versions of the pipeline where narrow features like atomic lines might be incorrectly identified as spikes (see Fig. 3 in Sloan et al. 2001). Because the occasional problem still remains, the *swsmake* code filters the corrected spectral data for spikes, as described in § B1.

We tried to avoid erroneously identifying narrow atomic lines as spikes while still removing all actual spikes from our data. In rare pathological cases our algorithm does not succeed. Our spike rejection algorithm relies on the assumption that a real line should appear in both scan directions, but a spike, even if it affects all detectors in a given scan direction, should not affect the other scan direction. In the case of HD 150193 (TDT 08200444, top panel of Fig. 15), spikes appear in most detectors in both scan directions at the same approximate wavelength (~19.0–19.5 μm), fooling the software into treating them as a real emission line and leaving them in the data. The opposite error occurs in one of the spectra of NGC 6543 (TDT 02400807), where the 10.52 μm [S IV] line exists in only one scan direction in the raw data used by the OLP 10.1 pipeline. Since it only appears in one scan direction, it mimics a spike, and the software removes it from the data (Fig. 15, *bottom panel*).

The relative importance of dark current and gain corrections can be complex in spectra with faint continuum levels and strong emission lines. While the correction to the emission line should be multiplicative, the correction to the continuum is additive, and our algorithm treats the entire segment as requiring an additive correction. For this reason, the ratios of line strengths between spectral segments should be studied cautiously in those cases where additive corrections have been applied. Since most emission line sources observed by SWS are extended, there are other reasons to be cautious about line strengths as discussed above.

A number of sources have emission lines in the overlap regions used to normalize the spectral segments. The most prominent of these lines is the Br α transition at 4.05 μm, which falls on the boundary between Bands 1E and 2A where the detector material, resolution, and noise characteristics change significantly. The two segments generally give different profiles and strengths for this line, which would adversely affect the normalization. For this reason the normalization process for Br α sources uses data from either side of the line, but not on it (see § B2). For some sources, such as Wolf-Rayet stars, the line is sufficiently broad as to have wings within the regions still used for normalization, and this may impact the final normalization in this region. A few sources have other lines in this overlap region, but these are ignored because of the small numbers involved.

In a few 1.NE and 2.E sources, the Humphreys series of hydrogen recombination lines occurs in the overlap region

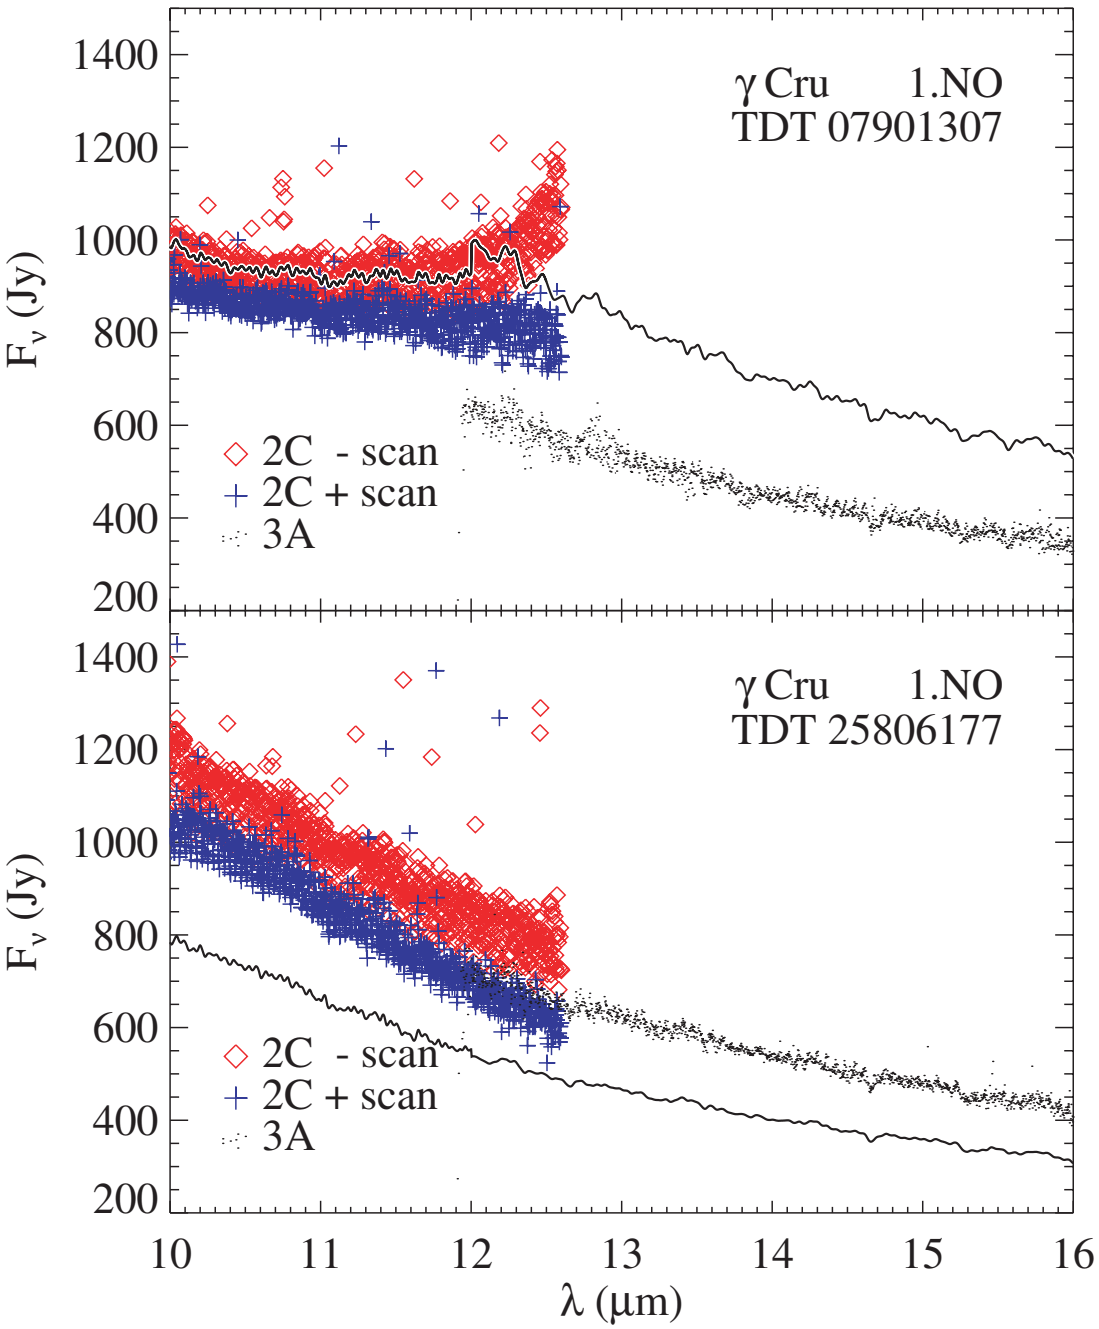

FIG. 14.—Two spectra of γ Cru, showing how the quality of the memory correction in Band 2C can affect the normalization to Band 3A. In the upper panel, *dynadark* has not completely corrected the memory in the dark current, producing a kink in the final normalized spectrum. In the lower panel, a better correction results in a more satisfactory interface between the two bands. Symbols are as in Fig. 2, with the addition of small dots for Band 3A.

between Bands 1D and 1E ~3.5 μm. It can be impossible to match both the continuum levels and the line strengths, as Figure 16 shows.

5.4. *Fringes*

Spectral fringing results from interference in internally reflected light in the filters and detectors. It is most pronounced in Band 3 because the vertical thickness of the thin BIB detectors is a few times larger than the wavelength. In general, the period of the fringes in wavelength space increases with wavelength. Also, the reflective back surface of the detector is not perfectly parallel to the front surface, causing the periodicity of the fringes to be a function of the position of the source within the aperture. Since the position of a source can vary during a scan, the period and amplitude of the fringes vary with wavelength in a complicated manner. The amplitude of the fringes will also depend on the angular extent of the source; point sources will produce larger fringes than extended sources. In general, fringing is more of a problem with higher-resolution data, and thus its effect on the spectra depends on the speed of the observation, as Figure 17 illustrates. Leech et al. (2002) and Kester, Beintema, & Lutz (2001) provide further information about fringes.

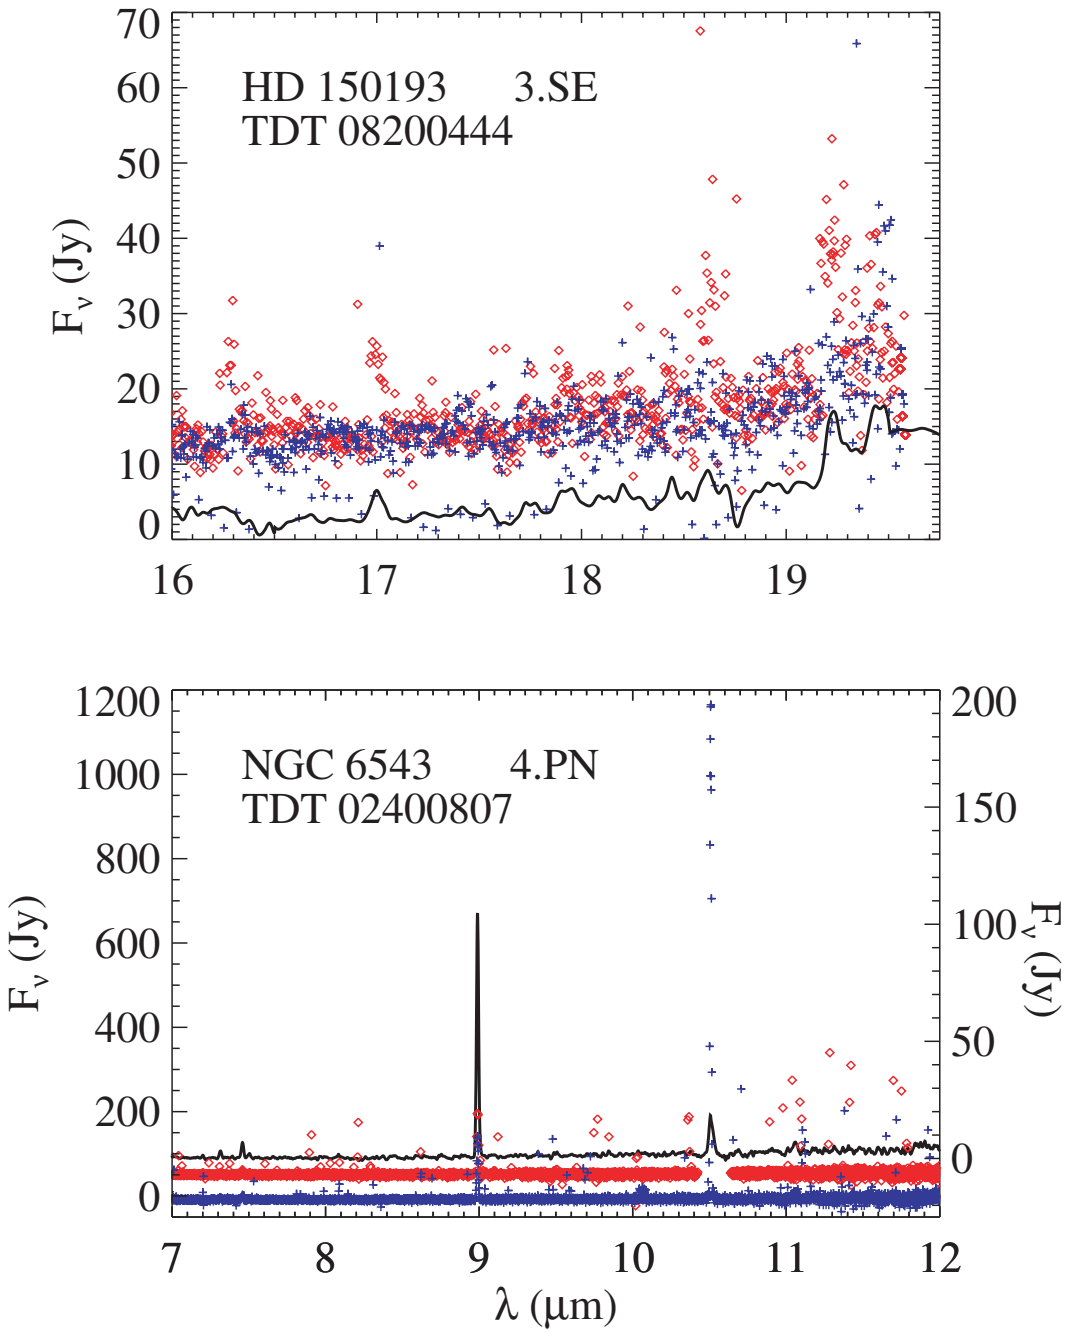

FIG. 15.—Examples of spikes in both scan directions surviving the rejection process (*top*) and an atomic line being rejected as a spike (*bottom*). Raw data in the two scan directions are plotted with red diamonds and blue plus signs (for " −" and " +" scan directions, respectively), while the final spectrum appears as a solid line. In the top panel, poor data between 19.2 and 19.6 μm in the spectrum of HD 150193 prevent the spike rejection algorithm from completely eliminating two spikes. This problem also affects the normalization of the neighboring band at 19.5 μm. In the bottom panel, the missing data at 10.5 μm in one scan direction leave the software with no way to distinguish the [S IV] line at 10.52 μm from a spike. In this panel, the two scan directions are offset vertically for clarity, and the final spectrum is plotted at a reduced scale from 0 to 200 Jy.

The defringing algorithms available (e.g., in OSIA or ISAP) are best applied interactively. As a consequence, our algorithm does not attempt to remove any fringes in the AOT 1 spectra. In the final spectra produced by *swsmake*, the fringe amplitude is typically 1%–2% of the flux (peak-to-peak) for speeds 1 and 2; for Speed 3, it is 2%–5%, and for Speed 4, 5%–10%. In general, the fringes appear only in Band 3 (most strongly in Band 3C). For fainter sources, the fringes are usually smaller than the noise. We estimate that the fringes appear at less than 2% strength in 86% of the *swsmake* database and at more than 5% strength in 5% of the database. For the AOT 6 spectra, the higher resolution of the spectra made defringing mandatory, so for these 23 spectra, fringes have been removed.

Many interesting absorption lines, particularly from molecules, exist at strengths hidden by the fringing in the AOT 1 data. Analysis of these types of features will require the user to begin with the AAR data or the more raw standard

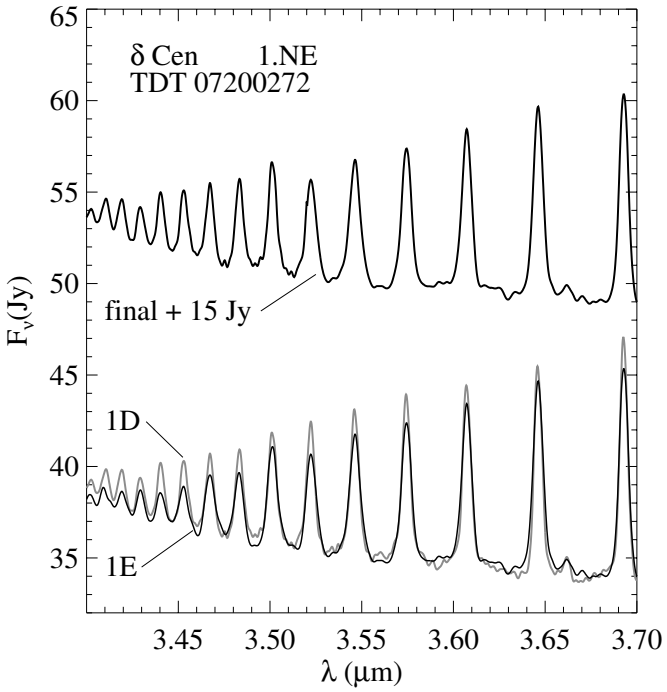

FIG. 16.—Overlap region between Bands 1D and 1E (segments 4 and 5), including the unnormalized ("pws") data for Bands 1D (*gray line*) and 1E (*solid line*), and the final ("sws") data (offset by 10 Jy). The continuum levels for 1D and 1E are well matched, but the line strengths differ. This leads to the discontinuity in line strengths at 3.52 μm, where the two bands are spliced together. Matching the line strengths instead of the continuum level would lead to a discontinuity in the continuum instead.

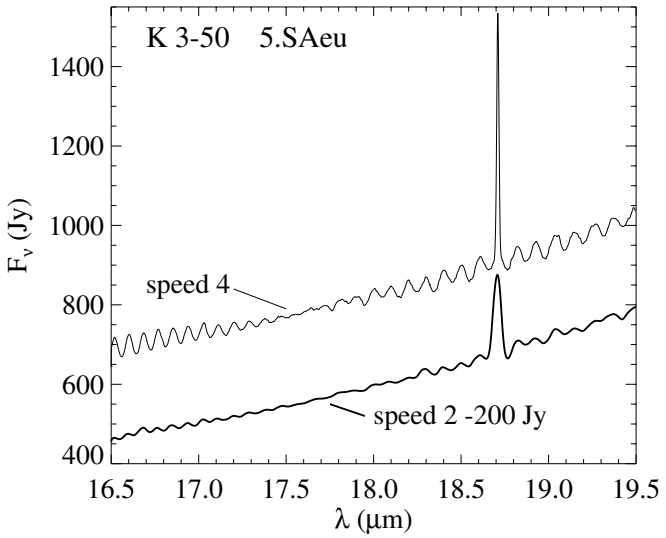

FIG. 17.—Band 3C data for K 3-50. Fringing in speed 4 data (*thin line*; higher resolution and sensitivity) and speed 2 data (*thick line*, offset by 200 Jy; lower resolution and sensitivity). The speed 4 data show more pronounced fringes. The speed 4 data are from TDT 38402466, and the speed 2 data are from TDT 14601350.

processed data (SPD) from *ISO* and perform the defringing by hand for AOT 1 spectra (preferably using OSIA).

5.5. *Spurious Features*

While the pipeline has been well tuned to eliminate spikes from the data, the occasional spurious feature can still appear. Particularly strong spikes which affect many detectors simultaneously can lead to residuals in the spectrum, as Figure 6 shows for α Vir at 16 and 22 μm. Band 4 is particularly prone to noise and spikes, especially at low flux levels. Because of this, care should be taken in interpreting putative features in Band 4 in sources where none are expected (in spectra classified in Groups 1 and 2 in particular).

As Figure 18 shows, residuals after the RSRF correction can also appear in the data. As discussed previously these tend to result when the additive corrections of the dark current are inaccurate, leading multiplicative corrections like those for the RSRF astray. These artifacts are most noticeable at 9.3, 10.1, 11.1, and 12.3 μm. We shifted the boundary between Bands 2C and 3A to 12.4 μm to avoid the 12.3 μm artifact in Band 3A.

In addition, spectral structure within Band 3E (27.3–27.7 μm) must be treated carefully, due to the notoriously poor behavior of Band 3E and the difficulty of aligning it with its neighbors. Figure 12 illustrates this problem, and § 4.3 describes it. The "pws" data provide a means of checking the validity of spectral structure in this region.

5.6. *Effect of Intrinsic Differences in the Spectra*

The SWS observed a large variety of objects with a wide range of spectral energy distributions (SEDs). As a result, some of the caveats described above will affect some spectral types more than others. Kraemer et al. (2001) have outlined some examples. For instance, stellar sources (Groups 1 and 2) have sharply falling SEDs and rarely have significant flux past ~25 μm. Cooler objects, such as H II regions and planetary nebulae (Groups 4 and 5), on the other hand, often have very low signal-to-noise for $\lambda \lesssim 12$ μm. Thus, the cooler objects will be more susceptible to uncorrected memory effects in Band 4 than the warmer sources.

The overall SED of a source primarily affects the reprocessing results by determining the starting point for the normalization process, either Band 1E for warmer sources or Band 3D for cooler sources.

The SWS database spans a wide range of brightnesses, and the brightness can also influence which calibration issues are significant. For bright and faint sources, the relative impact of dark current offset and gain uncertainties differ, and this affects how the spectra are processed. For fainter sources, the dark current dominates, so corrections (both within and between segments) are additive. For brighter sources, the gain uncertainties dominate, so corrections are multiplicative. There are other ramifications as well. As one example, the memory effects in Band 2 are less of a problem for fainter sources.

5.7. *Changes during the Mission*

During the *ISO* mission, the SWS instrument team made numerous changes to their methods of observation and calibration, improving the quality of the observations over time. To help the user of the database judge the impact of the problems discussed above, we outline below the modifications that affect the data and what those effects are.

The performance verification phase took place during orbits 0–79. Observations during these orbits were not necessarily in a standard mode, and they should be treated with more caution than observations taken later.

During the first ~200 orbits, strategies for the dark current measurements and photometric checks were adjusted.

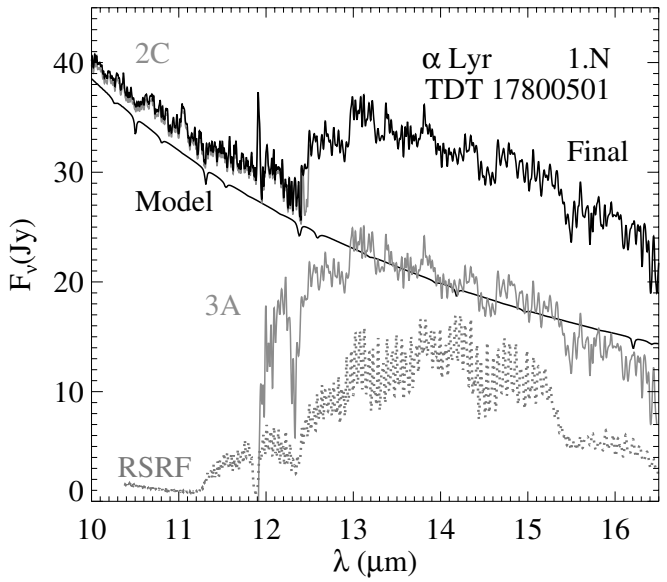

FIG. 18.—Comparison of the RSRF and a spectrum of α Lyr, showing how, in some cases, residuals from the RSRF can manifest themselves in the final spectrum. The strongest artifact at 12.3 μm clearly arises from the RSRF, and additional fainter features in Band 3A also appear to be artifacts.

Thus, for observations early in the mission, attempts to optimize these strategies sometimes led to additional problems with the data.[7]

For most of the early revolutions, the dark current was only measured prior to the spectral scan, as opposed to both before and afterward. If the dark current drifted during the scan, as often occurred, a single measurement could be a poor characterization of the dark current. This could lead to differences in the up and down directions, for instance, and incorrect spectral shapes. For high-speed observations (speeds 1 and 2), scan 3 had darks only prior to the scan until Revolution 123. For low-speed observations (speeds 3 and 4), all scans had only prior darks available until Revolution 191. After these orbits, measurements of the dark current were made both before and after each scan. The lack of a dark measurement at the end of scan 3 makes the *dynadark* correction for Band 2C less reliable than in later observations where darks before and after the scan help pin down the memory effects. Figure 14 illustrates this problem; the earlier spectrum in the top panel is less reliable due to the missing dark measurement.

Until Revolution 191, low-speed observations (speeds 3 and 4) also had photometric checks made before every scan in Bands 2, 3 and 4. These checks required illuminating the detectors with a bright internal radiation source, and due to the memory effects in Bands 2 and 4, this led to problematic dark current and flux measurements. The development of *dynadark* helps counter this problem in Band 2, but Band 4 remains irrecoverable for some sources (see Fig. 4 in Kraemer et al. 2001). Roughly 60 observations were made with this photometric pattern.

The pointing model also improved during the mission (Salama 1998). In Revolution 369, the star tracker was recalibrated, and in Revolution 451, a new Sun model was implemented. These changes affect the normalizations between segments.

At Revolution 64, the scan ranges changed. For most bands, the result is that the first 24 observations are normalized over slightly smaller ranges than for the rest of the mission (the "log" files include the reduced ranges). In addition, a gap appears in these spectra near 27.7 μm since Band 3D ended at $\sim$27.65 μm and Band 3E did not start until $\sim$27.75 μm. Since data in Band 3E are unreliable past $\sim$27.7 μm, the gap in the final "sws" spectrum extends to $\sim$28.03 μm, the start of Band 4 in these early observations. One AOT 6 spectrum (TDT 04800954) has a comparable gap in the data at $\sim$27.7 μm.

6. SUMMARY

We have created a database of 1239 AOT 1 spectra from the SWS, uniformly reduced and renormalized, along with 23 full-scan AOT 6 spectra. The spectra are publicly available on-line in a simple format consisting of a short header and the data in four-column ASCII format, as explained in Appendix A. A log file describes the processing steps and normalization constants for each observation. The IDL software which implement the algorithm described in Appendix B can be accessed through the same site, along with supporting documentation. We have assessed the quality of the spectra in this database (§ 4) and discussed the relevant caveats to be borne in mind by future users (§ 5). We feel that the overall quality of these data are sufficient to make them a valuable contribution to the astronomical community. We hope that they prove useful.

This work was supported in part by a NASA grant for the analysis of *ISO* dedicated-time observations. K. E. K. would like to thank the National Research Council for support via a Research Associateship through the Air Force Office of Scientific Research. G. C. S. wishes to thank Cornell for travel and computational support during the completion of this project. We thank Thijs deGraauw for support of the STARTYPE proposals and Harm Habing and Martin Kessler for contributing their assigned observing time to the project. We are grateful to E. Sturm for permission to reproduce his data. An anonymous referee helped us make substantial improvements to this manuscript by providing one of the most thorough reports we have encountered in our careers. The First Look software is Copyright 1995, 1996 California Institute of Technology, who acknowledge U.S. Government Sponsorship under NASA Contract NAS7-918. ISAP is a joint development of the Long Wavelength Spectrometer and SWS Instrument Teams and Data Centers. Contributing institutions to ISAP include the Centre d'Étude Spatiale des Rayonnements (CESR), the Institut d'Astrophysique Spatiale (IAS), the Infrared Processing and Analysis Center (IPAC), the Max Planck Institut für Extraterrestrische Physik (MPE), the Rutherford Appleton Laboratory (RAL), and the Space Research Organization of the Netherlands (SRON). OSIA is a joint development of the SWS consortium, including SRON, MPE, the Katholiede Universiteit Leuven (KUL), and the Astrophysics Division of the European Space Agency (ESA). This research has made use of the SIMBAD database, operated at the Centre de Donnees Astronomiques de Strasbourg in Strasbourg, France.

[7] Leech et al. (2002) mentions some of these changes, but omits others, such as the change in dark current measurement at Revolution 123.

APPENDIX A

ACCESS TO THE SPECTRAL ATLAS

As of the time of publication, the atlas is available at the *ISO* Data Archive (IDA)[8] or NASA's Infrared Science Archive.[9] Users will need to register at the IDA to access the data files there. Further information about the spectral atlas can be found either by contacting the first author directly or viewing his Web site.[10]

APPENDIX B

THE *swsmake* ALGORITHM

As outlined in § 2, the *swsmake* algorithm works in two stages. It first forms one spectrum from the 24 available for each spectral segment, and it then combines these together into a single continuous spectrum by normalizing the segments together and removing the regions of overlap. This Appendix presents a more complete description of the algorithm.

B1. FITTING A SEGMENT TO ITS MEDIAN

To produce one spectrum for each spectral segment from the 24 spectra produced by the 12 detectors in each scan direction, the *swsmake* code produces a median spectrum to which each individual spectrum will be shifted. It first regrids the individual spectra to a common wavelength grid and filters them for any spikes or atomic lines to prevent these features from affecting the corrections which will follow.

At each wavelength in the new grid, the median is found for both scan directions. These two scan directions are averaged for all segments, with two exceptions. First, only the second scan direction is used for the median of Band 4 (segment 13) to minimize the impact of the memory effect. (Fig. 2 shows the memory effect in Band 2.) Second, when the medians from the two scan directions disagree by more than 20%, the software will use data from what it judges to be the more reliable of the two scan directions over the affected wavelength range. The log file will include a comment that the routine in question (*segmentmrg*) is using variable weights to combine the medians for that segment.

The *swsmake* code generates a set of wavelength-dependent corrections for each detector and scan direction by comparing the regridded and filtered spectra to the combined median. It then fits these with a polynomial so that the corrections are a smooth function of wavelength. This preserves detailed spectral structure while correcting for the broad deviations from the median expected from the variations of the dark currents and gain among detectors, and for Bands 2 and 4, the memory effect. For spectral segments where all fluxes are greater than 15 Jy and the maximum is greater than 20 Jy, the software assumes that multiplicative constants are appropriate. In this case, the corrections are fitted to a cubic function of wavelength. For the fainter sources (those where any flux is below 15 Jy or all flux is below 20 Jy), dark current subtraction should dominate gain errors, so the correction is additive. Due to the poorer signal/noise ratio in these fainter data, the corrections are limited to a linear function of wavelength.

The software then corrects all 24 spectra in the segment, in their *original* wavelength grids, so that their general shape follows the median, generating a correction at each wavelength by using the polynomial fits and either making an additive or multiplicative correction as appropriate for that segment.

The next step is to remove the few spikes that survive the improved spike-rejection algorithm in the OLP 10.1 pipeline. These result primarily from hits by cosmic rays, and they occur far more frequently in Band 4 than in the other bands (Heras et al. 2001). A single detector may be affected, or several detectors in one scan direction could be hit simultaneously. A spike might cover a single wavelength element, or it might have a long, exponentially decaying tail (an effect most common in Band 4). The *swsmake* code identifies spikes at each wavelength by comparing data from individual detectors and the two scan directions using a median filter. A possible spike which occurs at the same wavelength in both scan directions is assumed to be a real feature; if it only occurs in one direction, the software flags it as a spike (see Fig. 15).

B2. NORMALIZING SEGMENTS TO EACH OTHER

As described in § 2.3 the normalization algorithm picks the brighter of Bands 1E or 3D (segments 4 or 11) as the starting point and relies on good data in the overlap regions to eliminate discontinuities between all bands from 1A to 3D. The decision between an additive or multiplicative correction follows the same criteria as used to combine spectra within a segment. Once the software has normalized Bands 1A to 3D to each other, it then normalizes Band 4 to Band 3D. Since these two bands do not overlap, extrapolation is required, as explained in the following paragraph. Band 3E cannot be used to bridge the gap between Band 3D and Band 4 due to its poor behavior. Once Band 4 is normalized, Band 3E is fitted to Band 4 using the modified method described below.

The details of the extrapolation between Bands 3D and 4 depend on shape of the spectrum in the 20–35 μm range, as determined by the classifications in Paper I. Sources classified as Group 1, 2, or 3 (naked stars, stars with dust shells, and warm

[8] Accessible from http://www.iso.vilspa.esa.es under the "Data Archive" header.
[9] Accessible at http://irsa.ipac.caltech.edu/data/SWSAtlas/SWSAtlas.html.
[10] See http://isc.astro.cornell.edu/~sloan/library/swsatlas.html.

dust shells, respectively) present a Rayleigh-Jeans tail at these wavelengths. The software fits a quadratic over the wavelength ranges 20.0–26.7 and 27.7–31.0 μm and determines the normalization by comparing these functions at 27.2 μm.

Band 3E also requires special care. The average of this band (over 27.3–27.7 μm) is fitted to the average of Bands 3D and 4 (averaged over the intervals 26.3–27.3 and 27.7–28.7 μm, respectively).

Approximately $\sim$60% of the full-scan spectra in the SWS database are treated as described above, but there are two exceptions which require deviations from the standard method.

First, several classes of spectra show spectral structure in the 20–35 μm region that require a linear extrapolation technique over shorter wavelength ranges to fit Band 4 to Band 3D. Carbon-rich sources showing an emission feature $\sim$26–30 μm are one example (see Fig. 19). For sources classified as 3.CR, 4.CR, and 4.CT, the extrapolation fits lines over the ranges 26.0–27.3 and 27.7–34.00 μm. The lines are compared at 27.7 μm. Red spectra in Groups 4 and 5 (spectra which peak beyond $\sim$20 and $\sim$40 μm, respectively), are the other example. The extrapolation fits lines over the wavelength ranges 23.5–27.3 and 27.7–31.0 μm since the standard quadratic fit would underestimate the strength of Band 4 relative to Band 3D. The normalization is determined at 27.7 μm. The *swsmake* code recognizes the standard extrapolation as normalization *type* 0. The modifications for cool carbon-rich sources and red spectra are recognized as *types* 1 and 2, respectively.

The second exception arises from strong emission lines in some spectra which can affect the interpolation between bands, especially between Bands 1E and 2A, as explained in § 5.3. For sources in Groups 1.NE, 2.E, 4.PN, 4.PU, 5.PN, 5.UE, we modify the normalization between segments by excluding wavelengths in the immediate vicinity of the Br α line at 4.05 μm. Other lines appear at other boundaries, but present less of a problem because both bands respond to the line in a more similar manner. The *swsmake* code identifies the modified interpolation by adding 10 to the value of *type*. Thus, *type* 10 corresponds to the normalization for a stellar or relatively blue spectrum with emission lines. Table 5 shows the relation between the classifications of Paper I and the normalization *types*.

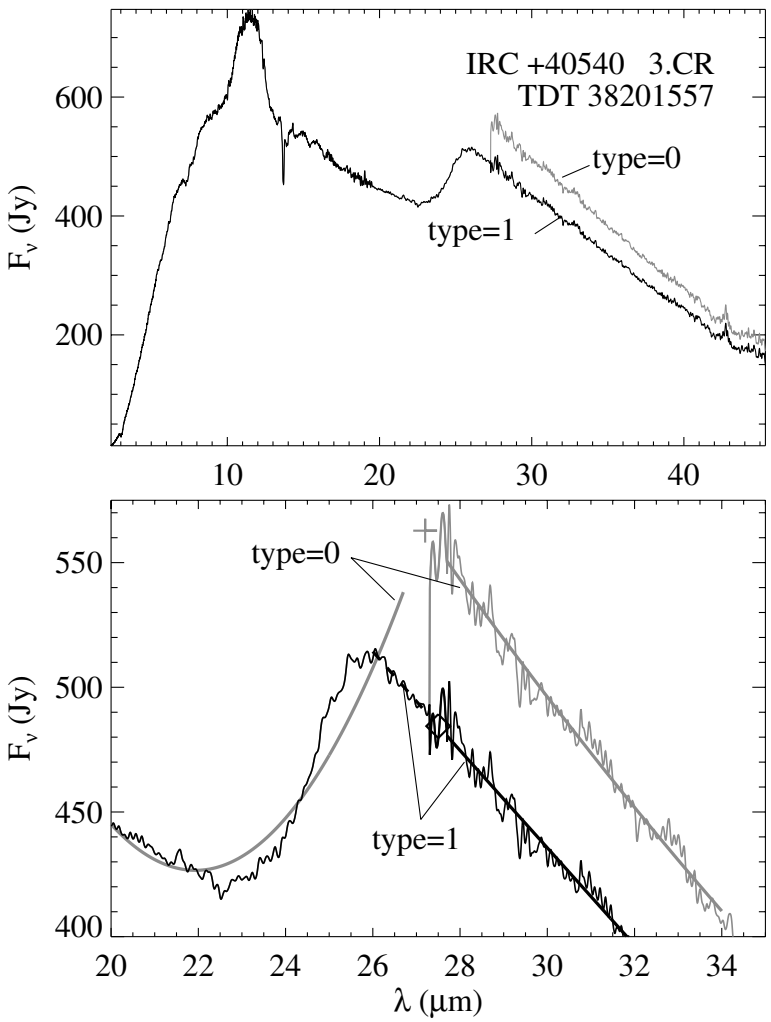

FIG. 19.—Comparison of different normalization methods for Bands 3D and 4, using the cool carbon-rich dust source IRC +40540 as an example. All data in both panels are shown *after* normalization. The top panel illustrates the results of the standard method of normalization (*type* = 0, gray) and the modified method used for carbon-rich sources with a 26–30 μm feature (*type* = 1, black). The bottom panel shows the details of the two normalizations. For the standard method (*type* = 0), heavy gray lines depict the polynomials fit to the data, and the "+" sign shows the point at which they are normalized. For the modified method (*type* = 1), the heavy black lines show the lines fit to the data (*dashed line* for Band 3D, *solid line* for Band 4), and the diamond shows where they are normalized.

TABLE 5
SEGMENT TO SEGMENT NORMALIZATION

Normalization *Type*	Infrared Spectral Classes
0	Group 1: N, NO, NC, NM
	Group 2: SEa, SEb, SEc, CE, C/SE, M, U
	Group 3: SE, SB, CE, W
	Group 7
10	1.NE, 2.E, 3.SAe
	Group 6
1	3.CR, 4.CR, 4.CN, 4.CT
11	4.PN, 4.PU
2	Group 4: SE, SB, SA, SC, SEC, C/SC, U/SC, F, M
	Group 5: SE, SA, U, F, M
12	Group 5: E, UE, PN

Table 6 gives the number of sources in each classification group which have been processed differently due to the presence of emission lines. While the table does not list the sources individually, the " log " file for a given observation always specifies the normalization method and the wavelength ranges used.

APPENDIX C

AOT 6 SPECTRA

The SWS obtained more than 800 spectra using AOT 6, and 23 of these covered the full 2.35–45.4 μm wavelength range available to the SWS (usually over a slightly reduced range, see Table 8). The SWS also observed 12 of these sources using AOT 1. As done in Paper I, the spectra were classified using the browse products initially, and where necessary, the results of our own processing. Table 7 gives the resulting classifications for all 23.

Of the sources also classified in Paper I, only the classification of Orion IRc2 differs, changing from 5.SAe (AOT 1) to 5.Eu (AOT 6). The difference results from the better sensitivity and higher resolution in AOT 6, which better reveal the emission lines and UIR features in contrast to the weak silicate absorption. This dependence on observing mode also occurred occasionally in the classifications of the AOT 1 spectra, where an observation taken at speed 3 or 4 might have a slightly different class than one taken at the lower-resolution speed 1 (e.g., OH17.7−2.2, classified as 4.SA at speed 3 but 4.F at speed 1). Also, the complicated emission structure in the vicinity of Orion IRc2 would make the spectrum dependent on pointing (the pointing of the telescope differed by 3″.5 for the two spectra, and long axis of the array had rotated by 15°). Spectra obtained in AOT 1 from nearby positions have classifications of 5.UE or 5.E.

The structure of the AOT 6 data resembled the AOT 1 data sufficiently to allow a modified version of the *swsmake* software to process the data available from the IDA. The differences in the structure of the data led to several modifications which are described below. The limited number of spectra enabled a more interactive approach where certain steps could be tailored for a given spectrum. In addition, the higher resolution made it necessary to remove the fringes from all bands, using the OSIA *fringes* routine with one modification: the low-wavelength limit for Band 3E was lowered from 27.5 to 26.0 μm. The log files include details about how the fringes were removed.

The AOT 1 observations consisted of one up-down pair of scans for each of the 12 spectral segments, making a total of 24 scans. The AOT 6 observations, on the other hand, usually consisted of either 45 or 73 scans, including reference scans, although in some cases there could be as many as 188 scans. For example, AOT 6 usually observed in Bands 1A, 1B, and 1E with one (nonconsecutive) up-down pair each and 0–3 reference scans, but Band 1D was often observed with three pairs (two covering ~2.95–3.5 μm, one covering ~3.3–3.6 μm, and three reference scans at ~3.04 μm). The modified *swsmake* software

TABLE 6
EMISSION LINE SOURCES AND NORMALIZATION

Infrared Spectral Class	Number of 'e' Suffixes	Normalization *Type*
3.SE	8	10
4.SA	2	12
4.SB	1	12
4.SE	9	12
4.SEC	6	12
4.F	1	12
4.M	4	12
5.SA	16	12
5.F	3	12

TABLE 7
CLASSIFICATION OF AOT 6 SPECTRA

Name	TDT	R.A. (J2000)	Decl. (J2000)	Group	In Paper I?	Comments
WZ Cas	24801029	00 01 15.71	+60 21 19.7	2.CE	N	
NGC 604	82802566	01 34 32.88	+30 47 01.6	6	N	G
S Per	42500605	02 22 51.71	+58 35 12.0	3.SE	Y	
AFGL 618	86301602	04 42 53.60	+36 06 53.0	4.CN	Y	
NGC 1984	26601410	05 27 40.27	−69 08 04.8	4.F	N	G
Orion IRc2	66002132	05 35 14.26	−05 22 31.6	5.Eu	Y	
30 Doradus 1	75002201	05 38 33.50	−69 06 27.1	5.E	N	G, Propn
β Pic	72501593	05 47 17.10	−51 03 59.7	2.M	N	
IRC +10216	19900101	09 47 57.30	+13 16 42.3	3.CR	Y	
SW Vir	24700418	13 14 04.50	−02 48 24.8	2.SEb	N	
IRAS 15194−5115	29700401	15 23 04.91	−51 25 59.0	3.CE	N	
V CrB	25502252	15 49 31.21	+39 34 17.8	2.CE	Y	
ZZ Her	28100117	16 28 38.38	+41 52 54.4	2.SEa	Y	
NGC 6334I	33100101	17 20 53.57	−35 47 01.2	5.SAeu	Y	
GC Sgr B2	28702002	17 47 21.80	−28 23 14.0	5.F	N	
M17	33201303	18 20 22.37	−16 11 25.3	5.Eu	N	
T Lyr	36100832	18 32 19.99	+36 59 55.5	2.CE	N	
NC 83	28604933	19 13 55.82	+54 12 13.5	7	N	
W51 IRS2	31901604	19 23 40.19	+14 31 07.1	5.UE	Y	
WR 147	04800954	20 36 43.60	+40 21 08.0	2.E	Y	
AFGL 2688	33800604	21 02 18.70	+36 41 37.0	4.CN	Y	
NGC 7027	33800505	21 07 01.50	+42 14 10.0	4.PU	Y	
TX Psc	37501937	23 46 23.57	+03 29 13.7	1.NC	Y	

NOTES.—Classifications and comments are as defined in Paper I. Those used here include: (G) extragalactic source, (Propn) Name from original observer. While 30 Doradus 1, and M17 do not have S01 spectra corresponding to the exact S06 location, nearby spectra have classifications of 5.E for 30 Dor, and 5.UE, 5.U, and 5.E for M17. Units of right ascension are hours, minutes, and seconds, and units of declination are degrees, arcminutes, and arcseconds.

discards the reference scans because they are usually much noisier than the normal scans, reducing the number of scans to 15–20 pairs in most cases. (The 188 scans of 04800954 reduced to 42 pairs while the 174 scans of 29700401 only reduced to 85 pairs.)

The SWS obtained up-down pairs consecutively in AOT 1, but not in AOT 6, requiring the creation of input files to tell the software which scans to consider as pairs (recall that the *swsmake* software compares the up and down scans to eliminate spikes and confirm spectral features). The software generated "pws" and "sws" files in the same manner as with the AOT 1 files, but the header information reflects the (arbitrary) number of segments in the spectrum instead of the standard 12 for AOT 1.

The AOT 6 data have little out-of-band data, reducing the overlap between neighboring segments from different bands and making the normalization more difficult and uncertain. The interfaces between Bands 3A and 3C, and 3C and 3D, often show dips, probably due to a lack of data, although the overall normalizations seem reasonable. These dips are in the raw data and so are not a product of our processing; their position at the ends of the scans and presence in different source types suggests

TABLE 8
WAVELENGTH GRIDS AND RANGES FOR AOT 6

	WAVELENGTH SPACINGS (nm)		WAVELENGTH RANGES (μm)	
BAND	AOT 6	AOT 1, Speed 4	AOT 6	AOT 1, Speed 4
1A	0.100	0.125	2.38–2.60	2.35–2.60
1B	0.200	0.167	2.60–3.02	2.60–3.02
1D	0.200	0.167	3.02–3.52	3.02–3.52
1E	0.300	0.250	3.52–4.03	3.52–4.08
2A	0.300	0.250	4.03–5.30	4.08–5.30
2B	0.625	0.500	5.30–7.02	5.30–7.00
2C	0.750	0.500	7.02–12.01	7.00–12.40
3A	0.750	0.833	12.01–16.50	12.40–16.50
3C	0.750	0.833	16.50–19.50	16.50–19.50
3D	1.250	1.667	19.50–26.00[a]	19.50–27.30
3E	1.250	1.667	26.00–29.00	27.30–27.70
4	2.000	2.500	29.00–45.20	27.70–45.39

[a] If band 3E is unreliable, the 3D/4 edge is at 26.9 μm.

that they are artifacts. The overlaps are larger between segments in the same band, so the normalizations are better constrained in these cases. After normalization, scans within the same band were averaged together. The exception is TDT 29700401, where the extremely short scans (e.g., 12 pairs in Band 1D instead of the typical three) necessitated normalizing and merging within a band with ISAP prior to applying the modified *swsmake* software.

As in the AOT 1 data, the interface region between Bands 3D, 3E, and 4 is complex. The earliest observations have no data from 27.6 to 28.4 μm, including 04800954. Band 3E is often unreliable. Band 4 often extends down to 3D, but may have residuals in the RSRF in strong sources, especially $\sim$28–29 μm. If Band 3E behaved reliably, the *swsmake* code normalized it to Band 3D and then normalized Band 4 to Band 3E, as opposed to the procedure in AOT 1, where Band 4 was first normalized to Band 3D and Band 3E was then normalized to both its neighbors. If Band 3E did not appear reliable, then the normalization followed the procedure as in AOT 1.

Table 8 gives the final spacing and ranges of the wavelengths for each band, along with the corresponding values for speed 4 data in AOT 1 for comparison. Normalization ranges vary depending on the observation, so Table 8 does not include them; they do appear in the log files.

REFERENCES

Cohen, M., Walker, R. G., Barlow, M. J., & Deacon, J. R. 1992, AJ, 104, 1650

Cohen, M., Witteborn, F. C., Bregman, J. D., Wooden, D. H., Salama, A., & Metcalfe, L. 1996, AJ, 112, 241

Fouks, B. I. 2001, in The Calibration Legacy of *ISO*, ed. L. Metcalfe & M. F. K. Kessler (ESA SP-481; Noordwijk: ESA), in press

Fouks, B. I., & Schubert, J. 1995, Proc. SPIE, 2475, 487

Heras, A. M., et al. 2002, A&A, 394, 539

Heras, A. M., et al. 2001, in The Calibration Legacy of *ISO*, ed. L. Metcalfe & M. F. K. Kessler (ESA SP-481; Noordwijk: ESA), in press

Kester, D., Beintema, D., & Lutz, D. 2001, in The Calibration Legacy of *ISO*, ed. L. Metcalfe & M. F. K. Kessler (ESA SP-481; Noordwijk: ESA), in press

Kester, D., Fouks, B., & Lahuis, F. 2001, in The Calibration Legacy of *ISO*, ed. L. Metcalfe & M. F. K. Kessler (ESA SP-481; Noordwijk: ESA), in press

Kraemer, K. E., Sloan, G. C., & Price, S. D. 2001, in The Calibration Legacy of *ISO*, ed. L. Metcalfe & M. F. K. Kessler (ESA SP-481; Noordwijk: ESA), in press

Kraemer, K. E., Sloan, G. C., Price, S. D., & Walker, H. J. 2002, ApJS, 140, 389 (Paper I)

Leech, K., et al. 2002, The *ISO* Handbook, Vol. V: SWS - The Short Wavelength Spectrometer, Version 2.0, ed. T. Müller & J. Blommaert (Noordwijk: ESA)

Onaka, T., de Jong, T., & Yamamura, I. 2002, A&A, 388, 573

Price, S. D., Sloan, G. C., & Kraemer, K. E. 2002, ApJ, 565, L55

Salama, A. 1998, SWS Beam Profiles and *ISO* Pointing (report available from *ISO*)

Salama, A., et al. 2001, in The Calibration Legacy of *ISO*, ed. L. Metcalfe & M. F. K. Kessler (ESA SP-481; Noordwijk: ESA), in press

Shipman, R. F., et al. 2001, in The Calibration Legacy of *ISO*, ed. L. Metcalfe & M. F. K. Kessler (ESA SP-481; Noordwijk: ESA), in press

Sloan, G. C., Kraemer, K. E., Goebel, J. H., & Price, S. D. 2003, ApJ, submitted

Sloan, G. C., Kraemer, K. E., & Price, S. D. 2001, in The Calibration Legacy of *ISO*, ed. L. Metcalfe & M. F. K. Kessler (ESA SP-481; Noordwijk: ESA), in press

Sloan, G. C., & Price, S. D. 1995, ApJ, 451, 758

———. 1998, ApJS, 119, 141

Sturm, E., Lutz, D., Tran, D., Feuchtgruber, H., Genzel, R., Kunze, D., Moorwood, A. F. M., & Thornley, M. D. 2000, A&A, 358, 481

Vandenbussche, B., et al. 2002, A&A, 390, 1033