

Source-dependent master sky images for the WFC3/IR grisms

G. Brammer, R. Ryan, N. Pirzkal

November 2, 2015

ABSTRACT

*Extraction of WFC3/IR slitless spectra depends on an accurate determination of the diffuse background light that is observed in all grism exposures. The two-dimensional structure in the background of WFC3/IR grism exposures is caused primarily by 1) overlapping grism spectral orders that are vignetted at different locations within the detector field of view and 2) the spectrum of the diffuse background source. To remove this structure, the **aXe** analysis software (and related pipelines) fit and subtract “master sky” images, which have been generated for each grism from on-orbit science exposures. In addition to the zodiacal light (i.e., reflected Solar continuum) background seen by WFC3/IR, a $1.083\ \mu\text{m}$ emission line in the Earth’s upper atmosphere often appears in exposures obtained while the spacecraft is outside of the earth’s shadow and its intensity varies on timescales of an orbit or even a single sample sequence. The single **aXe** master sky images available for each grism do not take this spectral variation into account. Following the methodology used to create the **aXe** master sky images, we generate images for each of the background components separately that together provide a more accurate description of the background observed in any given exposure. The images are made available for download¹ and we provide an algorithm for applying them to observed WFC3/IR data.*

¹http://www.stsci.edu/hst/wfc3/analysis/grism_obs/calibrations/wfc3_grism_master_sky.html

1. Introduction

Slitless spectroscopy observations with the WFC3/IR grisms are almost always background-limited as the exposures are often long (e.g., two 1500 s exposures in an orbit) and the grisms are effectively broad passband filters. Note that what we refer to hereafter as the “background” is rather the diffuse *foreground* local to the Solar system. In slitless spectroscopy, this diffuse foreground is distinguished from an effective foreground/background of spectra from (resolved) sources external to the Solar System or the Galaxy that can overlap on the detector with the spectrum of the source of interest. Each pixel on the detector sees background light from nearby areas of sky at different effective wavelengths and even different spectral orders as defined by the dispersion characteristics of the grism itself. The observed background at a given pixel therefore depends on the spectrum of the background source and the spectral orders that contribute to it. While the background spectrum can be assumed to be constant across the field, the various spectral orders are seen at different locations across the detector field of view. As the dispersion is predominantly oriented parallel to detector rows, the order overlap results in vertical “curtains” across the detector, particularly dark bands along the left and right edges where the +2nd and 0th orders drop out, respectively, and don’t contribute to the background. These features can have sharp edges (high spatial frequencies) that would require very high order functions to provide an analytical description of the background.

As an alternative to modeling the 2D background analytically, representative empirical background-only images can be constructed from archival science exposures. Kümmel et al. (2011) prepared such “master sky” images for the WFC3/IR G102 and G141 grisms from the first ~ 100 science images obtained after the instrument was commissioned. Master sky images are constructed by combining exposures normalized to a constant value after masking pixels associated with sources and pixels otherwise flagged as unreliable in the data-quality arrays (e.g., hot pixels). With a sufficient number of input images of different “scenes” and therefore differentially masked input pixels, a high-fidelity image of the master background can be constructed at all positions across the detector. The **axE** slitless spectroscopy analysis software (Kümmel et al. 2009) scales and subtracts the master background image from the science exposures.

Brammer et al. (2012) noted that G141 observations from the 3D-HST large program showed a diversity of structure in the background that was not adequately explained by a single master background image. For example, some images showed a much sharper edge on the dark band along the left edge of the detector, which also was darker compared to the rest of the image. They noted some correlation between the structure and intensity of the background, but that the correlation itself was not sufficient to explain the diversity. As an empirical solution, they generated four separate master background images that could

adequately explain much of the observed diversity in the 3D-HST observations.

We recently discovered that the diffuse WFC3/IR background arises from two dominant sources: zodiacal continuum from sunlight reflected by dust in the Solar System and an *emission line* at $1.083\ \mu\text{m}$ from metastable Helium in the upper sunlit Earth atmosphere (Brammer et al. 2014). The intensity of the zodiacal background depends on a target’s position relative to the sun, which varies over the course of the year but is effectively constant during an exposure or visit (see Pirzkal 2014). The He line, however, is excited in the sunlit upper atmosphere and is thus only seen when the spacecraft is outside of Earth’s shadow. Its intensity can increase or decrease strongly *within* an exposure depending on the position of the spacecraft during the orbit and on the geometry of the telescope pointing with respect to the Earth limb (Brammer et al. 2014).

The two components of the WFC3 IR background are spectrally distinct (i.e., smooth continuum vs. single emission line), and each alone will result in a distinct structure of the 2D background structure observed in grism exposures. **In this ISR we explore the hypothesis that the diversity of backgrounds seen by Brammer et al. (2012) and other users is the result of the superposition of these two components with varying relative intensities.** Following techniques such as that described by Kümmel et al. (2011), we generate master sky background images for the G102 and G141 grisms from an ensemble of *individual reads* of the detector where the two background components can be effectively isolated.

2. Analysis

2.1. Datasets

Table 1 summarized the WFC3/IR datasets used for this study, which were publicly available in the MAST archive as of June 2014. The G102 datasets include a variety of programs with either fairly deep integrations on single fields (e.g., 12927) or short visits on many different fields (13420). The G141 datasets are taken from the 3D-HST and WISP large programs, as well as GO-11600, upon which the 3D-HST orbit/visit structure was based. Building on the initial work by Kümmel et al. (2011), the full dataset now comprises 307 and 772 individual exposures in the G102 and G141 grisms, respectively. Breaking the exposures into individual MultiAccum reads (most commonly with exposure times 100 s for the SPARS100 read sequence), we consider 3671 and 8959 separate reads of the WFC3/IR detector in the G102 and G141 grisms, respectively.

Grism	Program IDs	N_{visit}	N_{exp}	N_{read}
G102	12203, 12590, 12896, 12927, 12946, 13420	54	307	3671
G141	11600, 12177, 12328, 12902, 13352, 13517	234	772	8959

Table 1: Summary of datasets used.

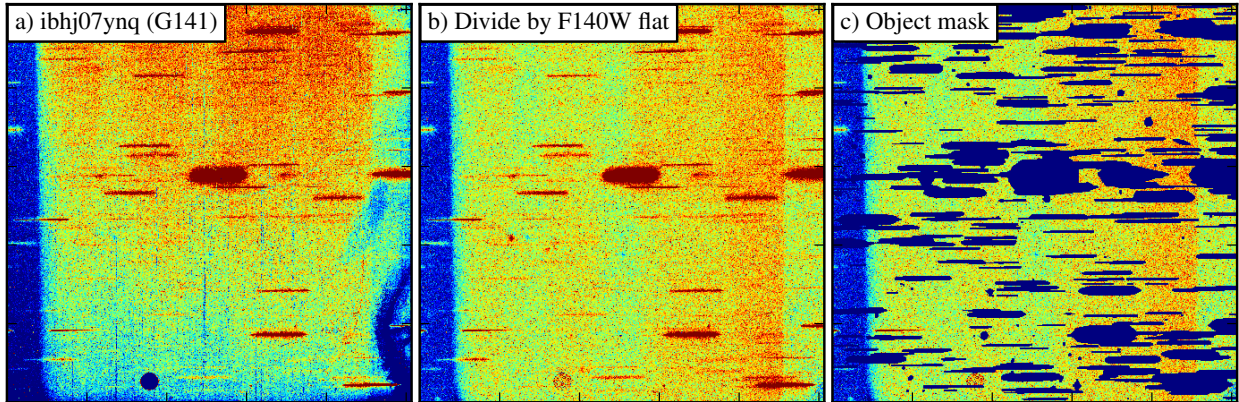


Fig. 1.— *a)* A normal FLT image produced by the CALWF3 pipeline. *b)* Dividing by the F140W imaging flat (uc721143i_pfl) separates the *multiplicative* flat-field features, such as the wagon-wheel at lower right, from the *additive* sky background. The ~ 100 pixel wide vertical bands at the left and right sides of the detector are the result of overlapping grism orders. *c)* Object spectra are masked in order to create master averaged sky images.

2.2. Object masking

We first create a mosaic of each visit combining its constituent exposures with the AstroDrizzle software (Gonzaga et al. 2012), which corrects for the geometric distortion of the camera and also flags hot pixels and cosmic rays not identified by the default *calwf3* pipeline processing. We detect “sources” on these rectified grism mosaic images using the SExtractor (Bertin & Arnouts 2002) software with an elongated kernel optimized to detect first-order grism spectra in the images (0th and higher orders with significant flux are identified as well). The SExtractor segmentation mask is expanded to further mask light in the extended outer regions of bright sources. The mask generated in the rectified frame is then “blotted” back to the distorted frame of each of the input exposures using the “*ablout*” routine to create an object mask for each exposure. Individual pixels identified as cosmic rays or otherwise having data-quality flags $DQ > 0$ are also added to the exposure mask. An example of an exposure mask is shown in Fig. 1.

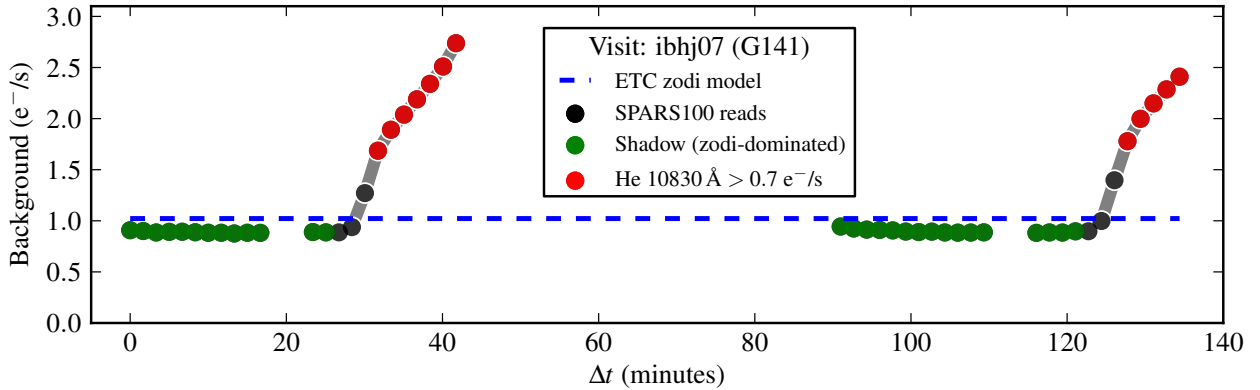


Fig. 2.— Background countrate for four G141 exposures from visit `ibhj07`, separated into individual 100 s reads from the SPARS100 sample sequence. Reads taken with the telescope in the Earth shadow are shown in green, where the background is dominated by zodiacal light. The dashed blue line shows the zodi prediction from the WFC3 ETC. Reads where the background from the He $1.083\ \mu\text{m}$ emission line exceeds $0.7\ \text{e}^-/\text{s}$ are shown in red.

2.3. Breaking up the ramps

A given WFC3/IR exposure consists of multiple non-destructive reads of the detector that are each recorded and sent to the ground. The accumulated charge is referred to as the exposure “ramp”, which is used by the `calwf3` pipeline to identify cosmic rays as large excursions from otherwise linear ramps. Observers typically use the pipeline-processed “FLT” products where the final count rate comes from a linear fit to reads “up-the-ramp”. Here we wish to consider each read separately as the relative intensities of the background components can vary considerably within an exposure, so we rather turn to the “IMA” products produced by the `calwf3` pipeline: for each read, i , the IMA files store the (calibrated) count rate r_i defined as the total *accumulated* charge c_i divided by the accumulated exposure time t_i since the reset read at the beginning of the sequence, i.e., $r_i = c_i/t_i$. To fully separate each read, we generate the *differential* count rate sequence from the IMA file that represents the average count rate within just that read, $\mathfrak{R}_i = (r_i t_i - r_{i-1} t_{i-1})/(t_i - t_{i-1})$. That is, \mathfrak{R}_i is the charge accumulated during read i divided by the sample time of that read. The datasets employed here use the SPARS50 and SPARS100 readout sequences with per-read sample times $(t_i - t_{i-1})$ of 50 and 100 s, respectively.

Fig. 2 shows differential count rate sequences for four exposures of a two-orbit G141 visit (`ibhj07` from GO-12177), demonstrating how the background count rate can vary significantly within a single visit/exposure. With knowledge of the spacecraft ephemeris

we identify individual reads obtained while the spacecraft was in the Earth shadow² and therefore when the background is expected to be dominated by the zodiacal component. As shown in Fig. 2, the ETC zodiacal light model (Dressel et al. 2014, §9.7.1; Pirzkal et al. 2014) provides a fairly accurate prediction of the zodi background intensity, accounting for the target coordinates and the epoch of observation (which therefore determine the heliocentric latitude and sun-target angle). We consider visits as groups and define the zodiacal level, \hat{Z} , averaged over the detector, for the visit as the median count rate of detector reads obtained with the spacecraft in shadow for that visit. We identify reads with a significant contribution from the He line as those with an observed background count rate $\geq 0.7 \text{ e}^-/\text{s}$ above the established zodi level. That is, $\hat{\mathfrak{R}}_i - \hat{Z} \geq 0.7 \text{ e}^-/\text{s}$, where the caret “^” denotes averaging over the unmasked pixels of the detector.

2.4. Constructing the master images

Before combining the individual read images into the master frames, we divide each frame by the broad-band imaging flat-field most appropriate for each grism³ (Fig. 1). This separates the multiplicative effects of the flat-field pixel-to-pixel variations from the additive backgrounds. The choice of flat fields to apply is not critical: in order to apply the master sky images derived here to science exposures, it is only necessary that *the same* flat-field image be divided into the science exposure before fitting and subtracting the background images. For use with `aXe`, the master sky images provided here should be multiplied by the flat-field files specified above in order to be able to fit them directly to the science exposures.

Each of the individual flat-corrected, zodi-dominated reads are normalized to unity median count rate and put into a large `masked array` with dimensions $N_{read} \times 1014 \times 1014$. For each pixel, after first rejecting any reads masked as object spectra we mask additional outliers beyond the (2.5, 97.5) percentile range over all available reads. The master zodi image is then created from the mean of all unmasked reads. A small fraction of pixels is masked in all available reads, either because they were all masked with object spectra or because they had static bad pixel flags set in the image data quality extensions. To avoid leaving empty holes in the final master sky images, we fill these empty pixels with the median of their 7×7 (non-empty) neighbors. For the master “Helium excess” images, we first subtract the master zodi image from each read scaled by the constant zodi level determined for its associated visit. The residual images are scaled to unity count rate⁴ and the master

²In practice, we determine the shadow timing from the SHADOENT and SHADOEXT header keywords in the JIF observation log files; see §7.

³We use F105W (`uc72113oi_pfl.fits`) for G102 and F140W (`uc721143i_pfl.fits`) for G141.

⁴Requiring a minimum $0.7 \text{ e}^-/\text{s}$ excess flux over the constant zodi level ensures that there will be sufficient

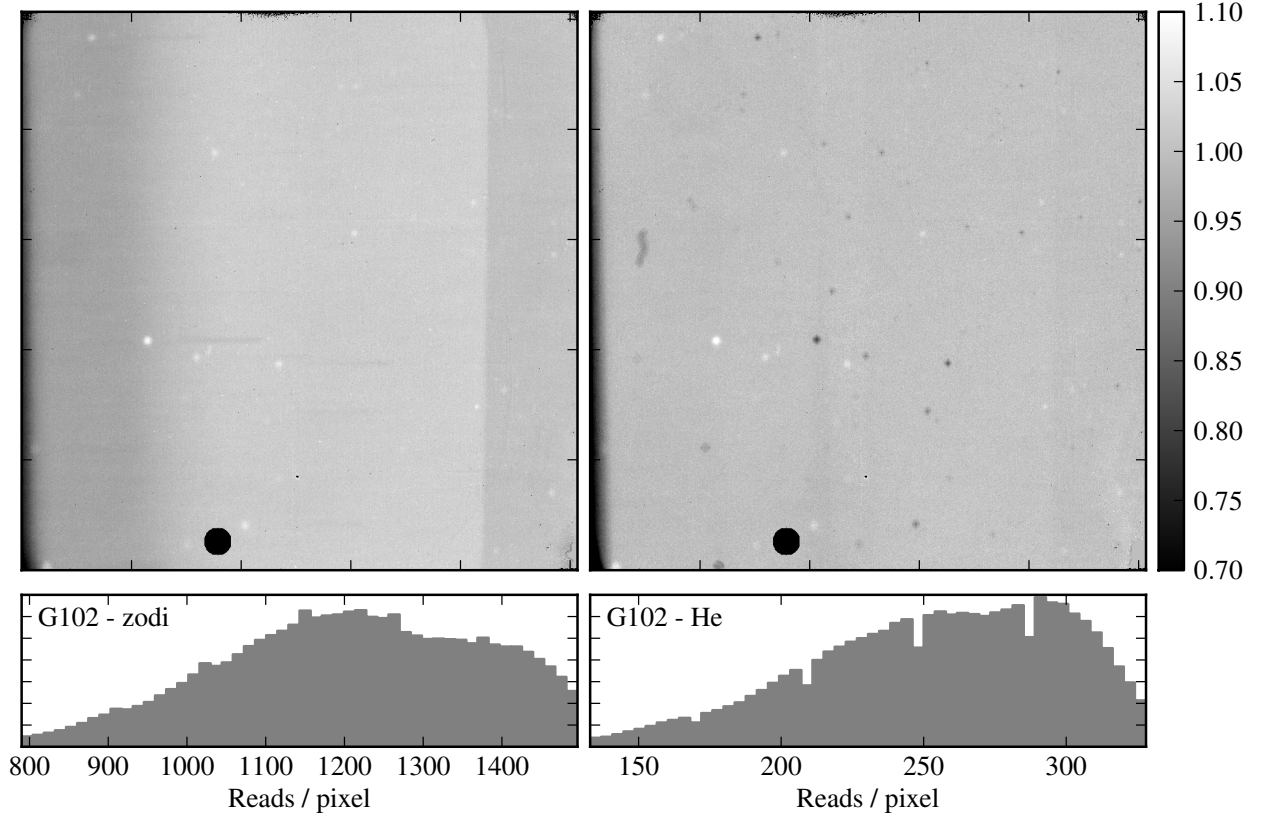


Fig. 3.— Master zodi and He background images for the G102 grism. The negative IR blob spectra clearly demonstrate the clean separation of the zodiacal continuum and He $1.083\ \mu\text{m}$ components of the background. Tick marks are shown every 200 pixels. The histograms in the lower panels show the distribution of the number of input reads that were averaged to make the final master image.

“excess” image is then constructed in the same way as described above for the zodi master image.

The master zodi and He excess images for both grisms are shown in Fig. 3 and Fig. 4. The differential background structure for the two components can be instantly appreciated: the shapes and contrast of the 100–200 pixel-wide bands on both the left and right sides of the detector are very different. The point-like positive excesses seen in both grisms result from the “IR blob” sensitivity depressions (Pirzkal 2010; McCullough et al. 2014) that become positive when dividing by the imaging flat-field. The grisms effectively disperse the blob, resulting in the negative spectra offset toward the right of the positive features.

residual excess flux left to scale and combine.




Fig. 4.— Same as Fig. 3 for the G141 grism. The spikes in the histogram are caused by aliasing in the histogram bins.

These are continuum spectra for the zodi component and emission line spectra for the He emission component, as expected (indeed these blob line spectra were used to identify the $1.083\ \mu\text{m}$ line as the background source). An interesting aspect of the G102 He image (right panel, Fig. 3) is the elongated vertical feature near $(x,y) = (100,600)$; this is a large IR blob just outside of the nominal imaging area and therefore not previously identified in the blob monitoring observations (e.g., McCullough et al. 2014).

Some residuals are apparent in the “wagon wheel” structure in the lower right corner of the G141 images in Fig. 4. This flat-field feature shows some wavelength dependence, so the residuals likely arise from the fact that we flat-field corrected the grism exposures with just the F140W imaging flat calibration image. It is worth noting that the flats themselves were created from science exposures in a similar way as the methodology described here (Pirzkal et al. 2011). For the F105W and F110W filters sensitive to the $1.083\ \mu\text{m}$ line, it may be worth creating flat-fields separately based on the background components as done here.




Fig. 5.— Object-masked, column-averaged G102 background profiles for individual SPARS100 reads dominated by zodiacal light (left) and for reads with a significant contribution from the He $1.083\ \mu\text{m}$ line (right). The rows are sorted following the observed background countrate of the reads, which is plotted next to the image panels. The zodiacal continuum and He emission-line backgrounds result in distinct background profiles, but once the components are separated, the background structure is very uniform across many exposures and programs.

The histograms in the bottom panels show the distribution of the number of acceptable reads that contributed to each pixel in the master image. For 50–100 s reads, the effective exposure time of the master images is 20–80 orbits for the zodi images and 5–10 orbits for the He excess images. It is worth noting that WFC3 GO programs are beginning to approach these depths (e.g., GO-13779, GO-14041) so the master sky images can not be treated as having infinite S/N but should rather be included in the error budget for the analysis of the deep pointings. Some systematic stripes remain in the G141 He image, most likely the result of imperfect masking of object spectra and a relatively few number of available reads. These issues may be addressed in a future analysis of additional data taken subsequent to the present analysis.




Fig. 6.— Same as Fig. 5 for the G141 grism. The G141 zodi profiles are quite uniform. However, some reads dominated by the He 1.083 μm line also show an additional component at the left part of the image from scattered light at low bright-limb angles.

2.5. Uniformity of the sky components

To explore the true separability of the sky components, in Fig. 5 and Fig. 6 we show column averages⁵ of all reads that go into creating the master images. Each row in these figures is the normalized column-average of a single read and the rows are sorted by the observed count rate of that read (left: $\hat{Z} \equiv \hat{\mathfrak{R}}_i$, right: $\hat{\mathfrak{R}}_i - \hat{Z}$), indicated in the curves next to the image panels. While the count rates themselves span a range of more than a factor of 5, the spatial structure seen in the backgrounds is remarkably consistent for the separated components independent of intensity. This is particularly true for the zodi component in both grisms. For the He component of the G141 grism, some non-uniformity is seen in the leftmost ~ 200 pixels and is likely caused by scattered (continuum) light from the bright earth limb observed at low target-to-limb angles (§6.10 in Rajan et al. 2010; see also §5). The worst cases of scattered light “blowout”, such as visit GO-11600 Visit 02, are excluded from the stacks, though some reads that remain in the stack could have low levels of scattered light not easily identified in the images. In any case, the structure in the G141 He reads at columns $x > 200$ is quite stable over the broad range of observed background intensity (0.7–2.5 e^-/s of the He component over zodi).

⁵The background structure seen in Fig. 3 and Fig. 4 is predominantly along detector columns.

3. Applying the background images to science exposures

Here we outline a strategy for using the master grism sky images developed here to subtract the sky from science grism exposures. The **aXe** software is not currently equipped to incorporate more than one background image so background subtraction with these images must be performed separate from and prior to **aXe** processing. **aXe** can then be run with the cleaned FLT images with the global background subtraction step turned off (`axeprep backgr="NO"`).

The first step of the background subtraction requires masking pixels with source spectra. In reality, very few if any pixels in the grism exposures will have *no* contribution from spectra of faint sources in the field. However, the background will dominate the flux for many pixels in typical exposures of sparse fields. Bright spectra can be identified with the SExtractor software as described above, with no prior knowledge of where spectra should fall based on the accompanying direct image. We provide a SExtractor convolution kernel useful for detecting spectra as objects at http://www.stsci.edu/hst/wfc3/analysis/grism_obs/calibrations/grism.conv. Using the direct image and the “FLUXCUBE” methodology, **aXe** can also be used to generate a full quantitative model of the grism exposures, which could then be used to generate the object mask. This approach requires running **aXe** multiple times to define the mask, subtract the background, and then extract spectra.

With the object mask in hand, the next step is computing the normalization factors necessary for scaling the master sky images to fit the science exposures, as the observed backgrounds will be a linear combination of the zodi and He background images presented here, perhaps with some additional component from scattered light. The scaling of the zodi component should be constant throughout all exposures of an orbit or multi-orbit visit and the He component can come and go depending on the characteristics of the spacecraft orbit during the visit. The normalization factors for the master sky component images can be determined with standard least squares techniques, for example, `scipy.optimize.leastsq` with an objective function that incorporates the constant zodi and variable He normalizations for the multiple exposures contained in a single visit. See also Appendix 6 for a complete example of an iterative linear inversion algorithm for determining the normalizations of the background components.

The results of this procedure for two G141 exposures from the visit shown in Fig. 2 are shown in Fig. 7 and Fig. 8. Here we show the masked and column averaged pixel values as a proxy for plotting the structure in the background. The blue curves show the result of using the single **aXe** image (Kümmel 2011), which was determined before the IR background components were well understood. While it provides a reasonable explanation for the zodi-dominated exposure at left, it results in very large residuals for the exposure at right that suffered dramatically increased backgrounds from the He component. The




Fig. 7.— Column-averages showing the observed and modeled backgrounds for two science exposures (the first two exposures shown in Fig. 2). The background in the exposure at left is dominated by the zodi component, while the exposure at right shows a significant contribution from the He component.

orange curves in Fig. 7 show the column profile for the model background image combined from the two master sky component images with normalizations computed as advocated above. The red curves show the best-fit sky model after including the additional scattered light component described in Appendix 5. For both exposures in Fig. 7, the combined sky-component background provides a much better description of the observed background than does any single component alone (e.g., the *aXe* master sky image).

Residuals from the background fits shown in the bottom panels of Fig. 7 are of order ≤ 0.01 $e^-/s/pix$, worse at the sharp edges of the “order-vignetting” structures at the edges of the detector but still only a small fraction of the background count rates themselves. Nevertheless, these residuals represent a significant systematic limitation on the ability to extract robust continuum spectra of faint sources: for comparison, a residual of 0.01 $e^-/s/pix$ is comparable to the continuum flux of a point source with magnitude $m = 23.7$ (AB) integrated over the brightest two central pixels of the trace. As the residuals are large scale features roughly parallel to detector columns, they can be reduced somewhat in a final step that subtracts the column-average residuals from the component-subtracted science exposures. This last correction is shown in the right panels of Fig. 8.

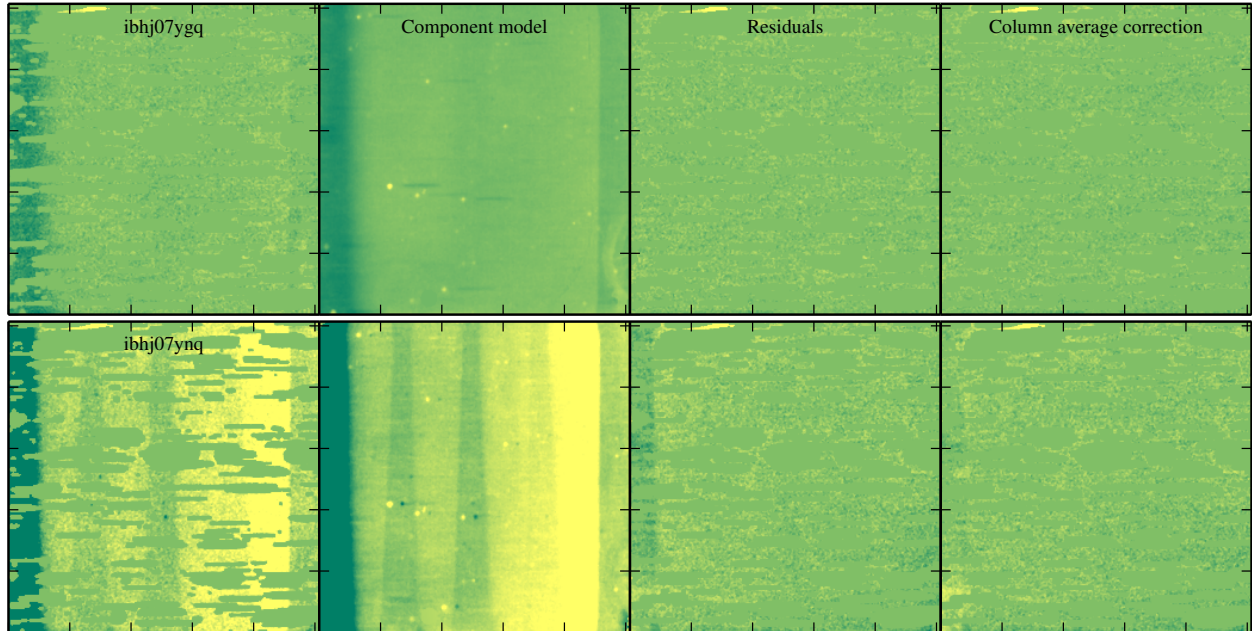


Fig. 8.— Full-frame images of two exposures showing the results of the background component modeling procedure outlined in §3. Object spectra were detected and masked with a spectrum-optimized SExtractor convolution kernel as described in the text.

4. Summary

We have created master sky images for the G102 and G141 WFC3/IR grisms separated into the two primary spectral components of the sky background, zodiacal light and $1.083\ \mu\text{m}$ line emission⁶. These master images provide a significant improvement in modeling and subtracting the global backgrounds observed in the IR grisms compared to results obtained using when only a single background image is used. Residuals from the background modeling are typically of order $\lesssim 0.01\ \text{e}^-/\text{s}$, 1% or less than the intensity of the background itself. Nevertheless, background modeling is likely a dominant source of systematic uncertainty for extracting continuum spectra of faint sources.

We would like to thank Susana Deustua for a careful reading of this manuscript.

References

Bertin, E. & Arnouts, S., 1996, A&AS, **117**, 393 [ADS]

⁶http://www.stsci.edu/hst/wfc3/analysis/grism_obs/calibrations/wfc3_grism_master_sky.html

- Brammer, G. B., et al., 2012, *ApJS*, **200**, 13 [\[ADS\]](#)
- Brammer, G. B., et al., 2014, WFC3 ISR 2014-03 “Time-varying Excess Earth-glow Backgrounds in the WFC3/IR Channel” [\[PDF\]](#)
- Dressel, L., et al., 2014, “Wide Field Camera 3 Instrument Handbook, Version 6.0” [\[URL\]](#)
- Gonzaga, S., et al., 2012, “The DrizzlePac Handbook” [\[URL\]](#)
- Kümmel, M., et al., 2009, *PASP*, **121**, 59 [\[ADS\]](#)
- Kümmel, M., et al., 2011, WFC3 ISR 2011-01, “Master Sky Images for the WFC3 G102 and G141 Grisms” [\[PDF\]](#)
- McCullough, P. R., et al., 2014, WFC3 ISR 2014-21, “Infrared Blobs: Time-dependent Flags” [\[PDF\]](#)
- Pirzkal, N., et al., 2010, WFC3 ISR 2010-06, “The WFC3 IR Blobs” [\[PDF\]](#)
- Pirzkal, N., et al., 2011, WFC3 ISR 2011-11, “Sky Flats: Generating Improved WFC3 IR Flat-fields” [\[PDF\]](#)
- Pirzkal, N., 2014, WFC3 ISR 2014-11 “The Near Infrared Sky Background” [\[PDF\]](#)
- Rajan, A., et al., 2010, “WFC3 Data Handbook, Version 2.1” [\[URL\]](#)

5. Appendix: scattered light

Some WFC3/IR exposures suffer scattered light that enters the instrument through an unintended path. This will be yet a separate component from the zodi and He background components studied above. It arises from scattered earth light and enters behind the filter wheel, so it is not effectively dispersed by the grism and does not show the same banding pattern caused by the vignetted grism orders. In some cases the scattered light is manifested as a “blowout” when the target pointing approaches the bright Earth limb and much of the left half of the detector shows increased flux (see the example in the right panel of Fig. 9). In other cases the increased flux might not be so severe to be easily seen but will be elevated over the nominal background light that enters the instrument through the intended light path. Considering reads with excess flux at the left edge of the detector such as those visible in Fig. 6, we fit another smooth component to help explain the structure coming from the scattered light. This image is shown in the left panel of Fig. 9 and can provide an improved characterization of the backgrounds when combined with the primary master sky components (see Fig. 7).




Fig. 9.— Scattered light component image. The reference image shown in the left panel is created from a high-order 2D polynomial fit to excess scattered light flux observed for a subset of the He excess exposures. An example of an exposure with scattered light “blowout” is shown in the right panel, where the zodi and He components have already been subtracted.

6. Appendix: Iterative Inversion

Here we present an alternative method for modeling the IR grism sky backgrounds that iteratively builds object masks and directly inverts the linear system. For the sake of notation, we indicate all images in a calligraphy font. As described above, we model the sky as single zodiacal component (\mathcal{Z}) for each visit and a separate helium emission (\mathcal{B}) for each image. Therefore the sky model (\mathcal{S}) is given as:

$$\mathcal{S}_{p,i} = z\mathcal{Z}_p + b_i\mathcal{B}_p \quad (1)$$

where p indicates pixel number (and serves to remind the reader that these are two-dimensional images) and i indicates the image number (and runs from 0 to $n - 1$, and n is the number of images). The unknown parameters are the amplitudes of the helium and zodiacal images: $\mathbf{x} = \{b, z\}$, respectively. Now the least-squares solution is found by minimizing:

$$\chi^2 = \sum_p \sum_i \mathcal{W}_{p,i} (\mathcal{F}_{p,i} - \mathcal{S}_{p,i})^2 \quad (2)$$

with respect to \mathbf{x} . Here, we have defined $\mathcal{W}_{p,i}$ as the weights for each pixel of each image (discussed below) and $\mathcal{F}_{p,i}$ as the observed images. This procedure results in a system of $n + 1$ linear equations of $\mathbf{v} = \mathbf{M}\mathbf{x}$. The elements of the matrix (\mathbf{M}) and vector (\mathbf{v}) are:

$$M_{i,i} = \sum_p \mathcal{W}_{p,i} \mathcal{B}_p \mathcal{B}_p \quad (3)$$

$$M_{i,n} = \sum_p \mathcal{W}_{p,i} \mathcal{Z}_p \mathcal{B}_p \quad (4)$$

$$M_{n,n} = \sum_p \sum_i \mathcal{W}_{p,i} \mathcal{Z}_p \mathcal{Z}_p \quad (5)$$

$$v_i = \sum_p \mathcal{W}_{p,i} \mathcal{B}_p \mathcal{F}_{p,i} \quad (6)$$

$$v_n = \sum_p \sum_i \mathcal{W}_{p,i} \mathcal{Z}_p \mathcal{F}_{p,i} \quad (7)$$

and of course \mathbf{M} is symmetric, so $M_{i,n} = M_{n,i}$. Now the least-squares solution is given by $\mathbf{x} = \mathbf{M}^{-1}\mathbf{v}$, however this formally allows for negative values of \mathbf{x} when the data are noisy. To avoid this unphysical solution, we solve for \mathbf{x} using standard non-negative least-squares methods. Now our algorithm for iteratively solving for sky background coefficients is:

1. Initialize the weights and sky models.

$$\mathcal{W}_{p,i} = \begin{cases} 0 & : \mathcal{D}_{p,i} \neq 0 \\ \mathcal{E}_{p,i}^{-2} & : \mathcal{D}_{p,i} = 0 \end{cases} \quad (8)$$

$$\mathcal{S}_{*,i} = \text{med}_p(\mathcal{F}_{p,i}) \quad (9)$$

where $\mathcal{E}_{p,i}$ and $\mathcal{D}_{p,i}$ are the observed uncertainty and data-quality images. At this stage, we assume the sky model for each image is a constant per image, and the next few iterations identify sources and determine the amplitudes of the spatial components.

2. Compute residuals and clip sources. The sky residuals are given by:

$$\mathcal{R}_{p,i} = \left| \frac{\mathcal{F}_{p,i} - \mathcal{S}_{p,i}}{\mathcal{E}_{p,i}} \right| \quad (10)$$

and then identify pixels with $\mathcal{R}_{p,i} \geq 3$ as containing sources and record the number of pixels (per image) which contain sources.

3. Grow source masks. We apply a binary dilation morphology operator to account for flux which is close to the residual threshold.
4. Update weights. We set the weights array to zero for any pixel now flagged as belonging to a source.
5. Solve linear system. As described above, we solve for \mathbf{x} using standard non-negative least-squares.
6. Update sky model. Using the new value of \mathbf{x} , we compute a new sky model following Equation 1.
7. Go to step 2. We repeat this iterative procedure until the number of pixels flagged as sources (per image) does not change. From experimentation, it takes typically five iterations.

7. Appendix: SHADOW timing

The code below demonstrates how to use header keywords in the “JIF” observation log files to determine which reads in a MultiAccum sequence were obtained with the spacecraft in the Earth shadow.

```

1 def in_shadow(file='ibhj07ynq_raw.fits'):
2     """
3     Compute which reads in a RAW file were obtained within the Earth SHADOW.
4
5     Requires the associated JIF files that contain this information, for example
6     "ibhj07040_jif.fits" for the default data file. These can be obtained by requesting
7     the "observation log" files from MAST.
8
9     """
10    import astropy.time
11    import astropy.io.fits as pyfits
12    import numpy as np
13
14    ##### Open the raw file
15    raw = pyfits.open(file)
16    NSAMP = raw[0].header['NSAMP']
17
18    ##### Find JIF file. Can either be association or single files
19    if raw[0].header['ASN_ID'] == 'NONE':
20        exp = raw[0].header['ROOTNAME']
21        jif = pyfits.open(exp[:-1]+'j_jif.fits')[1]
22    else:
23        exp = raw[0].header['ROOTNAME']
24        asn = raw[0].header['ASN_TAB']
25        jif = pyfits.open(asn.replace('asn', 'jif'))
26        for i in range(len(jif)-1):
27            if jif[i+1].header['EXPNAME'][:-1] == exp[:-1]:
28                jif = jif[i+1]
29                break
30
31    ##### Shadow timing (last entry and exit)
32    shadow_in = astropy.time.Time(jif.header['SHADOENT'].replace('.', ':'),
33                                  format='yday', in_subfmt='date_hms', scale='utc')
34
35    shadow_out = astropy.time.Time(jif.header['SHADOEXT'].replace('.', ':'),
36                                   format='yday', in_subfmt='date_hms', scale='utc')
37
38    ##### Array of read timings
39    t0 = []
40    for i in range(NSAMP):
41        h = raw['sci', i+1].header
42        ti = astropy.time.Time(h['ROUTTIME'], format='mjd', scale='utc')
43        t0.append(ti)
44
45    t0 = astropy.time.Time(t0)
46
47    ##### Test if reads were taken during shadow
48    test_in_shadow = ((t0-shadow_in).sec < (t0-shadow_out).sec) | ((t0-shadow_out).sec < 0)
49
50    return t0, test_in_shadow

```