Infrared Array Camera
Data Handbook

Version 3.0

January 20, 2006

Issued by the Spitzer Science Center

California Institute of Technology
Mail Code 314-6
1200 E. California Blvd.
Pasadena, California 91125 USA

http://ssc.spitzer.caltech.edu

help@spitzer.caltech.edu

Revision History

- Version 3.0: Jan 20, 2006

- Version 2.0: May 12, 2005

- Version 1.0: May 20, 2004

What is new in the Handbook?

- Jan 20, 2006: Included description of the flux calibration history, added text on the creation of uncertainty images, pipeline version S13 changes, included description of imask, added a description of scattered light rings, revised extended source calibration discussion, included description of superboresight pointing refinement. Additional small updates to text in many places.

- May 12, 2005: Revised description of the IRAC pipeline reflecting its current state, revised filenames, reorganized chapter on IRAC image artifacts and new sections in it, updated calibration information that reflects current status, new sections on array-location-dependent corrections, pixel phase corrections and pixel solid angle corrections, updated pointing information.

- May 20, 2004: Added a paragraph in the pointing section about HDR and short frame pointing header keywords.

The following people wrote the IRAC Data Handbook:

William T. Reach, Jason A. Surace, W. J. Glaccum, Sean J. Carey, Mark Lacy, Gillian Wilson, Patrick J. Lowrance, Brant Nelson, Pauline Barmby, Michael A. Pahre, Zhong Wang and Seppo Laine. This document was edited and compiled by Seppo Laine and Bill Reach.

Contents

Chapter 1

Introduction

This document was written for Infrared Array Camera (IRAC) data users to serve as a basic guide in the analysis of IRAC images. We have included a brief description of the IRAC data reduction pipeline that is used for the routine reduction and calibration of IRAC data. For further information, the user is referred to the IRAC Pipeline Description Document, available under the SSC/IRAC web pages. A PASP article describing the IRAC data reduction pipeline is also available under the IRAC web pages. We provide an overview of the data properties and a description of anomalies that the user may encounter in the IRAC data products. This Handbook reflects our current understanding of the instrument and the data, and is conceived to be an evolving document.

The Infrared Array Camera (IRAC) is an imaging camera operating at four wavelengths (called "channels"; 3.6, 4.5, 5.8, and 8 μm). The short wavelength detectors are InSb devices, while the two longer wavelength detectors are Si:As devices. The instrument observes two fields of view (FOVs) simultaneously, with their centers about $6.'5$ apart, leaving a gap of about $1.'5$ between the FOVs. Each FOV is seen simultaneously by two detectors (one InSb and one Si:As detector), fed through lenses and filters and through a beamsplitter. Each array has 256×256 pixels, with the same plate scale ($1.''2$ per pixel). The arrays are read out via four readout amplifiers. These readouts sweep across a row reading out all the array columns in the row, and then step to the next row. Fowler-sampling is used. At the beginning and end of each frame the array is read out non-destructively a set number (the "Fowler number") of times. The first set of readouts are accumulated negatively, then the last set positively, and the difference from frame start to end is generated. This accumulation in the digital signal processor boards is in a 24 bit format. The results are then bit-shifted by the Fowler number so that they fit into unsigned, 16 bit integers. Data are obtained simultaneously in all four arrays for all the frame times (except for the subarray 0.02 second frame time), regardless of whether only one or two FOVs were specified in the observing request, although the various channels may observe nearby, but separate, parts of the sky. The only exception is the stellar mode in which channels 1 and 2 obtain short exposures while channels 3 and 4 take a longer frame simultaneously.

There are two components to IRAC: the cold assembly and the warm electronics. The cold assembly primarily contains the detector arrays. It is temperature-controlled to milli-Kelvin levels and is contained within the cryogenically cooled multi-instrument chamber. The warm electronics is connected to the cold assembly by a cable.

In addition to this document, an IRAC data user is advised to consult the following documents, describing the instrument and the data:

- Fazio, G. G., et al. 2004, "The Infrared Array Camera (IRAC) for the Spitzer Space Telescope, ApJS, 154, 10

- *Spitzer Observer's Manual*, specifically chapter 6 describing IRAC

- Hora, J. L., et al. 2003, "Performance of the Infrared Array Camera (IRAC) for SIRTF During Instrument Integration and Test", SPIE Proceedings, Vol 4850, p. 83

- Reach, W. T., et al. 2005, "Absolute Calibration of the Infrared Array Camera on the Spitzer Space Telescope", PASP, 117, 978

Also, the IRAC Documents web page at

`http://ssc.spitzer.caltech.edu/irac/documents/`

lists a large number of other useful IRAC resources.

Chapter 2

IRAC Data Products

2.1 Description of Data Files

Table 2.1 lists the IRAC data files produced by the IRAC data reduction pipelines at SSC, together with brief descriptions of these files. The Basic Calibrated Data (BCD) are the calibrated, individual images. These are in array orientation and have a size of 256×256 pixels for the full array images, and 64 planes *times* 32×32 pixels for the subarray images. These data are fully calibrated and have detailed file headers. The Post-BCD pipeline combines the BCD images into mosaics (per wavelength and per frame time). Calibration observations designated as darks or flats go through a similar but separate pipeline that generates the products listed in the last section of Table 2.1.

Note that because of the "first frame effect" (see Section 3.2.1), the first frame of every Astronomical Observation Request (AOR) has a different delay time and it cannot be calibrated correctly. Therefore, the first frame of every AOR with a frame time greater than 2 seconds is taken in HDR mode which causes the first frame to be 0.6 seconds or 1.2 seconds in duration instead of the full frame time. This first frame usually has a name such as SPITZER_I1_11111111_0000_0000_2_bcd.fits and has no associated cosmic ray mask (brmsk file). The observer is encouraged not to use this first short frame. The pipeline mosaicker does not use it either when building the mosaic.

The BCD uncertainty files (listed below) are rough uncertainty estimations and do not include all of the systematic effects associated with IRAC detectors, nor do they include the absolute flux uncertainty. These uncertainty images are generated as follows. They begin as an estimate of the read noise (one number in electrons for the whole image) and the shot-noise due to the sky (proportional to the square-root of the number of electrons in the image). Then each module propagates the uncertainty image forward, including the uncertainties in dark and flat calibration files. The pipeline modules use the uncertainty image as a way to quantitatively estimate the quality of the sky estimate given by the value of a pixel. In the end, the uncertainty images overestimate the formal uncertainty of the image, because the net propagated uncertainty is much higher than the observed pixel-to-pixel fluctuations in the images. We therefore recommend that the uncertainty images only be used for relative weights between pixels, for example when performing outlier rejection or making a weighted mosaic that combines multiple input frames that view the same sky mosaic pixel. In the future we will provide a prescription for estimating the noise in an image that is validated with the in-flight data.

Please note that all of the calibration products (specifically, skydarks, skyflats, and linearity curves) are in the raw reference frame. Hence, the subarrays are located at pixel coordinates 9–40,9–40. In the BCDs, the subarrays are located in pixel coordinates 9–40, 9–40 in channels 3 and 4, and in pixel coordinates 9–40, 217–248 in channels 1 and 2. Also note that those

4

coordinates are for the case in which the first pixel is indexed as 1,1 (i.e. IRAF). E.g. IDL pixel indexes start from 0,0.

The pmask and dmask files are discussed in Section 4.2.1. The rmask files are discussed in Section 7.3.3.

2.2   IRAC Specific Header Keywords

Here we describe some of the important header keywords. A complete IRAC image header description is included in Appendix B.

AORLABEL is the name of the AOR as defined by the user in Spot when requesting the observations. The P.I. will be listed as the OBSRVR of each project. AORKEY (also known as REQKEY) is a unique digit sequence for each observation; it is also part of the filename for each BCD. EXPID is an exposure counter incremented within a given AOR for each data-taking command. Most data-taking commands generate multiple files: one per array in full array mode. The DCENUM is a counter of individual frames (per wavelength) from an individual command; it can be used to separate frames generated with internal repeats. The only observations with non-zero DCENUM currently are channel 4 BCDs for 100 second frame times. In high dynamic range mode, the long and short exposures are generated with independent commands and have different EXPIDs. Thus, for example, data from 12-second high dynamic range observations can be separated into long and short frames using the 'odd' or 'even' EXPIDs.

DATE_OBS is the time at the start of the AOR. Other times in the header include the time since IRAC was turned on (both for the beginning and end of the frame). FRAMTIME is the duration of the frame including Fowler sampling, and EXPTIME is the effective integration time. UTCS_OBS is the start of IRAC data taking sequence. ATIMEEND is the correct time of an integration end. HDRMODE tells you if the frame was taken in the high dynamic range mode.

BUNIT gives the units (MJy/sr) of the images. For reference, 1 MJy/sr = 10^{-17} erg s^{-1} cm^{-2} Hz^{-1} sr^{-1}. FLUXCONV is the calibration factor derived from standard star observations; its units are (MJy/sr)/(DN/s). The raw files are in "data numbers" (DN). To convert from MJy/sr back to DN, divide by FLUXCONV and multiply by EXPTIME. To convert DN to electrons, multiply by GAIN.

The predicted background (using the same model as accessed with Spot, evaluated for the wavelength, date, and coordinates of observation) is contained in three keywords: ZODY_EST, ISM_EST, CIB_EST. These are not based on the actual data from *Spitzer*. SKYDRKZB is the zodiacal background prediction for the skydark that was subtracted from the science image in the reduction pipeline. Thus the predicted background in the BCD data is ZODY_EST − SKYDRKZB.

DS_IDENT is a journal identification number for the Astrophysics Data System (ADS) to keep track of papers published from these data.

Absolute pointing information is contained in the following keywords. ORIG_RA and ORIG_DEC give the coordinates of the image center constructed from the telemetry using the Boresight Pointing History File, as indicated by the Boolean keyword USEDBPHF, and the file is listed in BPHFNAME. When pointing telemetry is not available, due to a telemetry outage, the commanded positions are inserted instead, USEDBPHF is false, and the coordinates will be less certain. RARFND and DECRFND are the refined positions derived by matching the brightest sources in the image with the 2MASS catalog. PA_RFND is the refined position angle of the +y axis of the image, measured east from north (CROTA2 is the same position angle but measured west from north). CRVAL1 and CRVAL2 give the coordinates of the image center, derived from the refined positions in all channels, and are usually the most accurate coordinates available.

Table 2.1: Sample IRAC file names.

Filename	Brief description
Basic Calibrated Data (BCD)	
SPITZER_I1_0008845056_0031_0000_01_dce.fits	Raw data
SPITZER_I1_8845056_0031_0000_1_bcd.fits	BCD data
SPITZER_I1_8845056_0031_0000_1_bcd.log	BCD pipeline log
SPITZER_I1_8845056_0031_0000_1_bunc.fits	BCD uncertainty file
SPITZER_I1_8845056_0031_0000_1_bdmsk.fits	BCD Dmask file
SPITZER_I1_8845056_0031_0000_1_bimsk.fits	BCD Imask file
SPITZER_I1_8845056_0031_0000_1_brmsk.fits	Rmask images
SPITZER_I1_8845056_0031_0000_1_ptn.log	BCD Pointing log
Post-BCD Processing	
SPITZER_I1_8845056_0000_1_E123458_msaic.fits	mosaic
SPITZER_I1_8845056_0000_1_A2987651_msunc.fits	mosaic uncertainty file
SPITZER_I1_8845056_0000_1_A2987653_mscov.fits	mosaic coverage file
SPITZER_I1_8845056_0000_1_A2987654_maicm.fits	HDR intermediate frame time mosaic
SPITZER_I1_8845056_0000_1_A2987655_muncm.fits	HDR intermediate frame time mosaic uncertainty file
SPITZER_I1_8845056_0000_1_A2987656_mcovm.fits	HDR intermediate frame time mosaic coverage file
SPITZER_I1_8845056_0000_1_A2987657_maics.fits	HDR short frame time mosaic
SPITZER_I1_8845056_0000_1_A2987658_muncs.fits	HDR short frame time mosaic uncertainty file
SPITZER_I1_8845056_0000_1_A2987659_mcovs.fits	HDR short frame time mosaic coverage file
SPITZER_I1_8845056_0000_1_E123456_reptg.tbl	list of DCE center positions in a mosaic
SPITZER_I1_8845056_0000_1_A987660_irsa.tbl	list of 2MASS sources used for pointing refinement
SPITZER_I1_8845056_0000_1_A2987662_mosnl.nl	Parameter file for mosaic/source extraction pipeline
Calibration pipeline data files	
SPITZER_I1_13450853_0000_1_C925234_sdark.fits	skydark
SPITZER_I1_13450853_0000_1_A2106548_scmsk.fits	skydark mask file
HDR_30s_1.2sf1d1r1_ch3_v1.2.0_dark.fits	IRAC labdark image (long frame)
HDR_30s_30sf16d1r1_ch3_v1.2.0_dark.fits	IRAC labdark image (long frame)
irac_b[1-4]_fa_superskyflat_091004.fits	IRAC superskyflat image
irac_b[1-4]_fa_20020921_lincal.fits	linearization calibration image
irac_b[1-4]_fa_20020921_lincal_cmask.fits	linearization calibration mask
irac_b[1-4]_fa_cdelt12_distort.tbl	array distortion table
irac_b[1-4]_fa_muxbleed_coeff_112003	muxbleed correction coefficients
irac_b[1-4]_fa_muxbleed_lut_100102	muxbleed correction look-up table
irac_b[1-4]_fa_pmask.fits	pmask
irac_b[1-4]_fa_pmask_flip.fits	flipped pmask (not in use)
irac_b[1-4]_fa_slmodel_v1.0.0.fits	subtracted scattered light model
irac_b[1-4]_fluxconv.tbl	flux conversion file used
irac_b[1-4]_mosaicprf.fits	PRF used in mosaicing
irac_b[1-4]_PRF.tbl	table of each PRF
irac_b[1-4]_PRF_1,2,3,4.fits	PRFs for each channel
README-1,2,3	archive log for calibration files

The BCD +x-axis (bottom, or horizontal axis) is in the direction of the telescope +Y-axis, and the BCD +y-axis (left side or vertical axis) is in the direction of the telescope +Z-axis.

Next we give an example of how an AOR file translates into final data products. A *Spitzer* observation is specified by a small list of parameters that are listed in the ".aor" file. This file is generated when the observation is designed (using Spot). The ".aor" file for a planned or performed observation can be retrieved using the "view program" feature of Spot. (You will be prompted for the program name or ID, which can be obtained from the image header keywords PROGTITLE and PROGID.) Here is an example AOR file:

```
#  Please edit this file with care to maintain the
#  correct format so that SPOT can still read it.
#  Generated by SPOT on:  5/9/2003      12:10:9

HEADER: FILE_VERSION=7.0, STATUS = PROPOSAL

      AOT_TYPE:  IRAC Mapping
     AOR_LABEL:  IRAC-FLS-CVZ-a
    AOR_STATUS:  new

 MOVING_TARGET:  NO
   TARGET_TYPE:  FIXED SINGLE
   TARGET_NAME:  FLS-CVZ
  COORD_SYSTEM:  Equatorial  J2000
      POSITION:  RA_LON=17h13m05.00s,  DEC_LAT=+59d10m52.0s
OBJECT_AVOIDANCE:  EARTH = YES, OTHERS = YES

        READOUT_MODE: FULL_ARRAY
              ARRAY: 3.6_5.8u=YES, 4.5_8.0u=YES
          HI_DYNAMIC: NO
          FRAME_TIME: 12.0
      DITHER_PATTERN: TYPE=Cycling, N_POSITION=5, START_POINT=1
        DITHER_SCALE: small
 N_FRAMES_PER_POINTING: 1
MAP: TYPE=RECTANGULAR, ROWS=7, COLS=6, ROW_STEP=277.0, COL_STEP=280.0,
     ORIENT=ARRAY, ROW_OFFSET=0.0,COL_OFFSET=440.0,N_CYCLE=1
SPECIAL: IMPACT = none, LATE_EPHEMERIS = NO,SECOND_LOOK = NO
RESOURCE_EST: TOTAL_DURATION=5848.4, SLEW_TIME=1089.0, SETTLE_TIME=1045.0,
SLEW_OVERHEAD=180.0, SPECIAL_OVERHEAD=0.0, UPLINK_VOLUME=9026,
DOWNLINK_VOLUME=59065440, VERSION=S7.0.B2
INTEGRATION_TIME: IRAC_3_6=60.0,IRAC_4_5=60.0,IRAC_5_8=60.0,IRAC_8_0=60.0
```

For this AOR, there are 210 files (6 columns × 7 rows × 5 dither positions) of each type for each channel. The final data products from this AOR in channel 2, provided it got assigned the AORKEY 6213376, are as follows:

```
SPITZER_I2_0006213376_0000_0000_01_dce.fits
SPITZER_I2_6213376_0000_0000_1_bcd.fits
SPITZER_I2_6213376_0000_0000_1_bcd.log
SPITZER_I2_6213376_0000_0000_1_bunc.fits
```

```
SPITZER_I2_6213376_0000_0000_1_bdmsk.fits
SPITZER_I2_6213376_0000_0000_1_bimsk.fits
SPITZER_I2_6213376_0000_0000_1_brmsk.fits
SPITZER_I2_6213376_0000_0000_1_ptn.log
....
SPITZER_I2_0006213376_0209_0000_01_dce.fits
SPITZER_I2_6213376_0209_0000_1_bcd.fits
SPITZER_I2_6213376_0209_0000_1_bcd.log
SPITZER_I2_6213376_0209_0000_1_bunc.fits
SPITZER_I2_6213376_0209_0000_1_bdmsk.fits
SPITZER_I2_6213376_0209_0000_1_bimsk.fits
SPITZER_I2_6213376_0209_0000_1_brmsk.fits
SPITZER_I2_6213376_0209_0000_1_ptn.log
```

After the name of the telescope, the first partition gives the instrument ("I" = IRAC), and the number after the "I" gives the channel (in this case, 2). The next part gives the AORKEY, then we have the EXPID, DCENUM, and the version number (how many times these data have been processed through the pipeline). One should generally use only the data from the highest version number, in case multiple versions have been downloaded from the archive. To verify that the data are from the latest pipeline version, check the CREATOR keyword in the header. Finally, there is a group of letters that specify what kind of data are in the file (see Table 2.1 above), and the file type (usually "fits" or "log"). The post-BCD file names include telescope name (SPITZER), "I" (for "IRAC"), the channel number, the productid (not the same as the AORKEY), the DCENUM, the (pipeline) version, "ensemble product id," the type of the data and the suffix. In the case of an ensemble product, "DCENUM" in the filename refers to the first DCE that was used in the ensemble creation, and "version" refers to the version of that first DCE. The letter "C" stands for "calibration" product: in the case of a calibration product, "DCENUM" refers to the first DCE that was used in the calibration creation, (pipeline) "version" refers to the version of that first DCE, and number after the "C" letter is the "calibration number". Note that for a given AORKEY of science data being retrieved, the AORKEY for the associated calibration products is different.

A list of 2MASS sources for the field of the IRAC observation is included in the data delivery as *irsa.tbl. Note that the 2MASS magnitudes given in the *irsa.tbl file are not meant for scientific use. For scientific use of the 2MASS data, query the 2MASS catalog directly from IRSA, and take into account the flux quality flags.

Chapter 3

IRAC Science BCD Pipeline

3.1 Overview

The IRAC BCD (Basic Calibrated Data) pipeline is designed to take a single raw image from a single IRAC detector and produce a flux-calibrated image which has had all well-understood instrumental signatures removed. The following section describes several of the major modules of the basic data reduction pipeline for science data. Similar pipelines are used for reducing calibration data. More detailed information on the pipeline can be found in the *Pipeline Description Document* (PDD).

The IRAC pipeline consists of two principal parts: the data reduction modules and the calibration server. Individual modules are written as stand-alone code and each corrects a single instrumental signature. Each module uses FITS files and text configuration files as input and produces one or more FITS files and log files as output. The calibration data, such as the skydarks and the skyflats needed to reduce a given image, are produced via a calibration server. A raw IRAC image is thus "passed" between successive modules, and at each step it becomes closer and closer to a finished, fully reduced image.

3.2 Significant Pipeline Modules

3.2.1 Dark Subtraction

The true dark current in the IRAC detectors is actually very low — the most notable dark current features are the electronic glows seen in the Si:As arrays (channels 3 & 4). However, the IRAC arrays experience considerable pedestal offsets which are commonly of the order of tens of DN. These offsets are dependent on the Fowler sampling, exposure time, and operation history of the arrays, and are believed to be due to very small thermal changes in the internal IRAC cold electronics. The most significant of these offsets is the "first-frame" effect: laboratory measurements show that the dark patterns and DC levels change as a function of the time elapsed between the end of the previous frame and the start of the current frame (called "delay time"). The first frame of a series of exposures is most affected, and therefore this effect is called "the first-frame effect." Figure 3.2 shows how the DC levels of darks change as a function of delay-time.

Due to the decision not to use the shutter in IRAC for dark and flat measurements, we have a somewhat sophisticated dark subtraction procedure. There will be two steps for the dark subtraction, one using a dark from the ground-based laboratory measurements (the so-called "lab dark"), and another using a delta dark which is the difference between the lab dark and the sky dark measured at a low zodiacal background light region.

Dark Subtraction I: First-Frame Effect Correction

In the first step, we subtract a calibrated lab dark from the data. Dark frames in all four IRAC channels are not constant. A dark image taken with a particular frame time and Fowler number depends on the amount of time elapsed since the previous image, and the frame time and Fowler number of the previous image. In a sequence of images with the same frame time and Fowler number, the variation from one image to the next consists principally of a uniform change over the whole image. To a much smaller extent, there is also a change in the relative offsets of each of the four outputs of an array (pin-striping) and a small spatial gradient across the image. This gradient can be striking when mosaicing several dithered images within an AOR. True dark current and multiplexer glow carry with them the usual shot noise in the number of charges collected at the integrating node, but the rest of the Fowler bias presumably is due to relaxation in the multiplexer. The noise in pixels with low dark current and low glow is the same as the read noise from the multiplexer, so the dark offset variations do not add to the pixel noise.

In general, the first frame of any sequence of images tends to have a different offset from the others, because it tends to have an interval, frame time, or Fowler number different from its preceding sequence (different AOR). If a frame is taken within a very short interval from the preceding image, (i.e., when multiple frames are commanded at once using repeats) the labdark calibration is smaller, but changing in value rapidly (Fig 3.2) and the calibration will be worse. Most observers should be dithering between frames, meaning that the interval between frames will be in the range 8 to 40 seconds where the labdarks are best calibrated.

The first-frame effect was calibrated in ground tests with the shutter closed. Figure 3.2 shows the dependence of dark offset on the interval between 30 second frames. Each symbol is the mean dark signal in one output of each array, averaged over a box near the center of the array. The dashed curves are multiple exponential fits to the means in the set of 48 images, with the first image excluded. The lower portion of each panel shows the residuals, in electrons, between the data and the fitted curves. The residuals are small compared to the expected sky backgrounds in 30 second frames. Channel 2 has very small variations in offset, and channel 4 has significant variations, though they are very small compared to the background. The most significant first-frame effect by far is in channel 3. The lab dark calibration images are derived from multiple expansion fits to individual pixels.

In the current pipeline version S13, the delay time, or the time since the previous IRAC image, is calculated from the header keyword `ATIMEND` of the previous frame and the current one. The module selects a lab dark of the appropriate Fowler number, frame time, and delay time from a table of possible lab darks. It then subtracts the calibrated lab dark from the image. After the first frame, a lab dark of the correct frame time and Fowler number can be used on all other frames until another dither or slew causes the time between frames to change. The delay time of the first frame of an observing sequence is often greater than 60 seconds (due to the target acquisition slew) so a standard labdark is used causing these first frames to have a slightly different calibration. *Therefore, the first frame of every AOR with frame time >2 seconds is taken with the high-dynamic-range mode to mitigate this relaxation time between the first frames. The observer is encouraged not to use the first short frame, and it will not be used in the pipeline mosaicker.*

Dark Subtraction II: Skydarks

The applied labdarks do not account for changes in the bias on-orbit or for any possible temporal variations in the bias. To calibrate these variations, skydarks are applied to all IRAC BCDs. IRAC will always be exposed to the sky, and hence all IRAC images will contain flux from the diffuse celestial background (such as zodiacal light) as well as compact objects such as stars and galaxies. As a result, the second part of the dark subtraction method strongly resembles
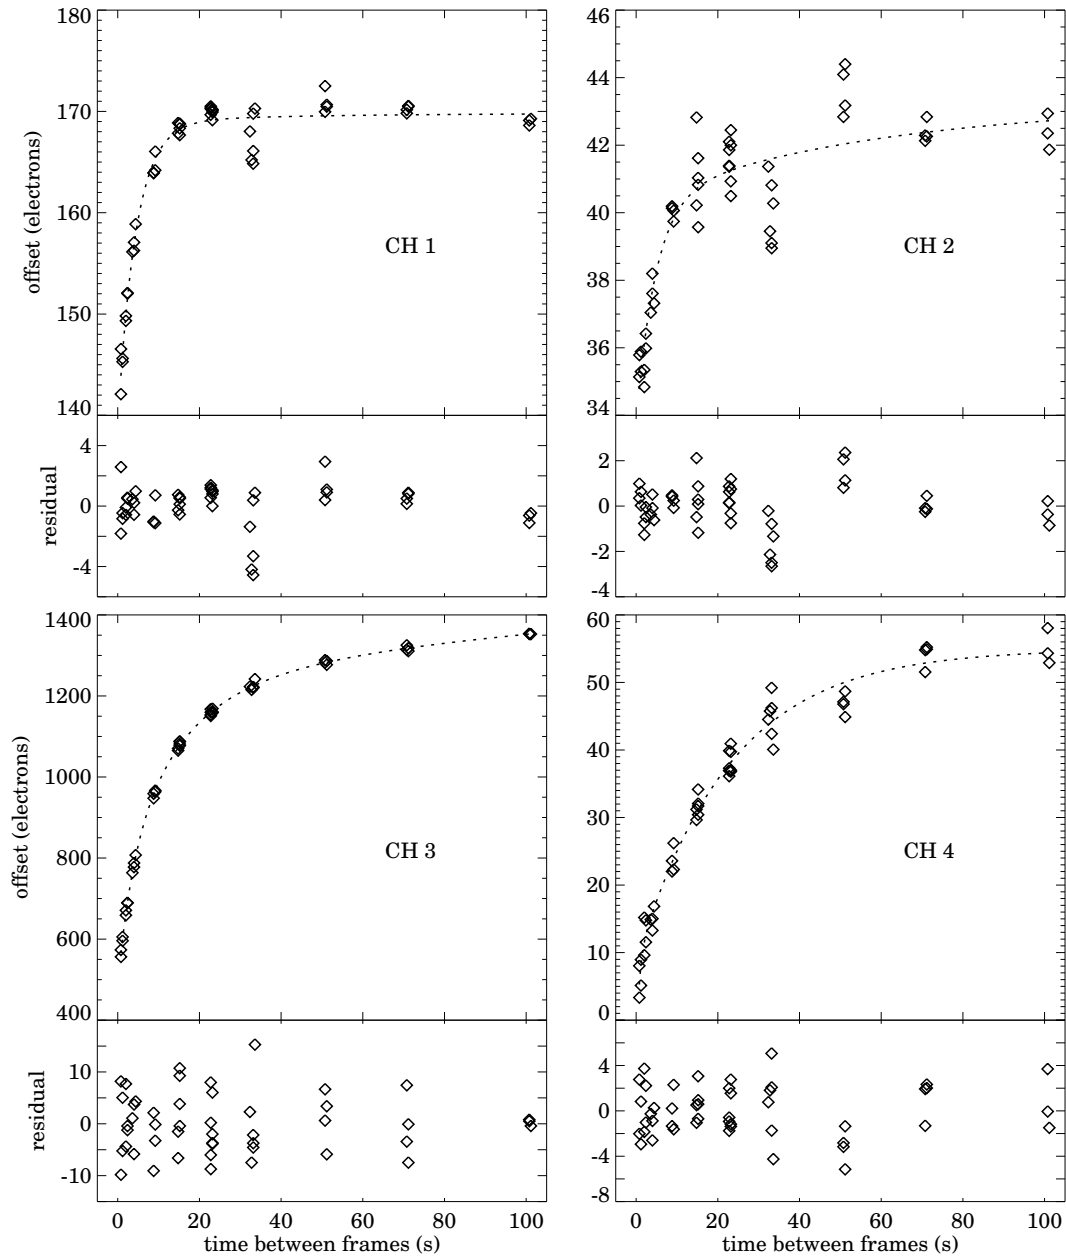
Figure 3.2: Illustration of the "first-frame effect." Each data point gives the mean level of 30-second frames taken at the specified interval (x-axis) from a previous image. The dashed line shows a fit and the inset panels show the residuals from this fit.

11

traditional ground-based data reduction techniques for infrared data.

At least twice during each campaign, skydarks of all Fowler numbers and frame times are observed. The skydarks are taken in pre-determined regions which have been specifically chosen to be the darkest parts of the sky and are in regions that are relatively free of bright objects such as stars.
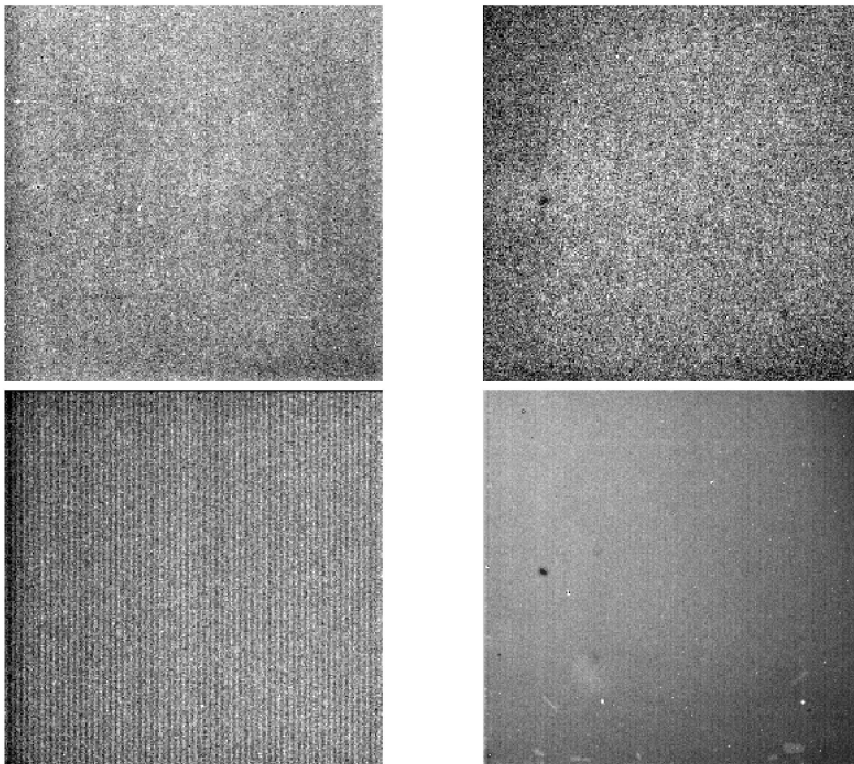

Figure 3.3: IRAC skydarks with a frametime of 12 seconds. Dark spots in channels 2 and 4 are contaminants on a pickoff mirror, appearing in all images, dropping the flux on the order of 30%. These data were taken from program pid = 1108, AORKEY = 9552128.

These sky darks are subtracted from all the data later in the pipeline than the lab darks. Since the sky darks contain any additive terms from the bias or dark current, these are subtracted out from the data. This is nearly identical to the "median sky subtraction" in ground-based data. Because this naturally over-subtracts the real data, we include an estimate of the sky brightness (based on a model) in the header (ZODI_EST) of the darks.

Each IRAC array is read out through four separate channels, and the pixels read out are arranged vertically, repeating every four columns. Small drifts in the bias level of each readout, particularly relative to the calibration labdarks, can produce a vertical striping called the "jail-bar" effect. This is most noticeable in very low background conditions. Before pipeline version S13, this was corrected in all channels by averaging the mean effect and adjusting individual channels by a common mean offset. In on-orbit tests, the mean offset and correction was found to be negligible, except in channel 3 data. Therefore, for S13 pipeline reprocessing and later versions, the darkdrift correction will only be applied to channel 3 data. The derived correction values for each channel are located in the header in the keywords DDCORR1, DDCORR2, DDCORR3, and DDCORR4. The overall background term determined is DDBKGND.
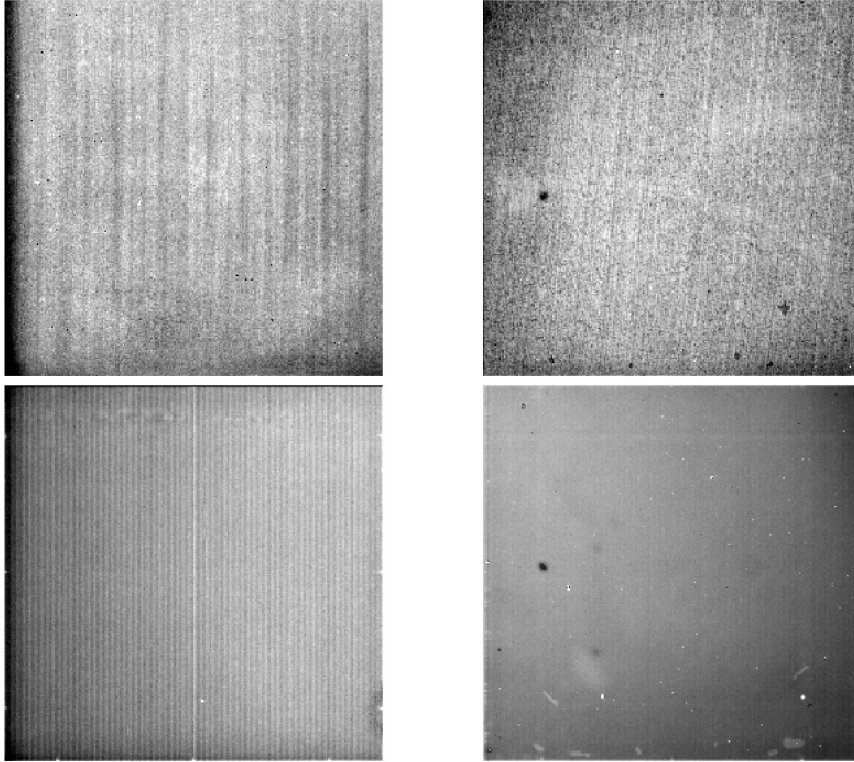
Figure 3.4: IRAC skydarks with a frametime of 100 seconds. The bright region in the corner of the channel 3 array (Si:As) is due to the electronics near the edge of the array. Dark spots in channels 2 and 4 are contaminants, appearing in all images, dropping the flux on the order of 30%. These data were taken from program pid = 1108, AORKEY = 9553152.

3.2.2   Multiplexer Bleed

The InSb arrays suffer from an effect known as "muxbleed." This is believed to be a result of operating the arrays at unusually cold temperatures. When a bright source is read out, the readout multiplexers (muxs) in the cold electronics do not return to their quiescent state for a considerable length of time. The result is a ghosting along the pixel readout channels, sometimes referred to as "ant trails." The effect is easily noticeable against a low background (such as a dark current measurement), and can extend the full length of the array. The muxbleed flux is not real — it is not "borrowed" from the actual source and as such needs to be subtracted, unlike CTE smearing in CCDs.

It appears that a pixel bleeds only as a result of the light falling onto it, and not as the sum of the value of the pixel plus the bleeding from previous pixels. Since we know the readout order of the pixels, we can start by correcting all pixels downstream from the first pixel, and then move on to the next pixel. The generalized solution is to correct this effect by a lookup table. The exact shape of the curve and the characteristic decay time vary as a function of incident intensity, so the correction is scaled by the intensity of the bleeding pixel according to

$$DN'_{n+4i} = DN_{n+4i} - LUT(i)/f(x) \tag{3.1}$$

where

$$f(x) = A(E/DN'_n) + B(E/DN'_n)^2 + C + D(E/DN'_n)^3 \tag{3.2}$$

and LUT is an entry in a lookup table for position n, and A, B, C, D, and E are scale factors.

13

Both the scale factors and the LUT are fixed for all pixels in a given array. Pixels with values below 5000 DN are not corrected, because the values in these cases will be just a few times the read noise. The muxbleed is not corrected in the subarray observations.

Observers should note that their calibration darks are not muxbleed-corrected. Muxbleed occurs in these images due to the presence of hot pixels. However, this occurs equally both in the darks and in the science frames and has been found to subtract noiselessly from the science data. Thus, the dark frame muxbleed is best simply considered a feature of the darks.

The pipeline muxbleed correction is far from perfect as the nature of the correction has changed in-flight. Work on improving the muxbleed correction is ongoing. More information about muxbleed, including sample images, can be found in Section 4.3.2.

3.2.3 Detector Linearization

The IRAC arrays have a non-linear response near full-well capacity, such that fluxes at levels above 30,000 DN (in the raw data) will be low by several percent. As part of pipeline processing, the data are linearized based on ground calibrations (which have been verified in flight) of this effect. The linearization is based on a quadratic model for all channels.

$$DN_{obs} = DN - \frac{A}{m^2}DN \qquad (3.3)$$

The arrays are assumed linear at zero DN, and coefficients m, A and C are based on ground-testing. The arrays saturate at values of 30,000, 35,000, 45,000, and 50,000 DN for channels 1, 2, 3, and 4. The linearization that is performed will mark "non-linear" pixels that exceed the DN value above which the linearization solution can no longer meet its required accuracy of 1% by setting a bit in the dmask. Because the data are Fowler-sampled and only differences between pedestal and signal reads are returned, once the data have saturated, an increasing exposure will actually *decrease* the returned DN as opposed to reaching and staying at a fixed high value. This is a result of saturating the signal reads, but not the pedestal reads. Because of this, the non-linearity mask bit will not be set for all saturated data.

3.2.4 Flat-Fielding

Similarly to all imaging detectors, each of the IRAC pixels has an individual response function (i.e. DN/incident photon conversion). To account for this pixel-to-pixel responsivity variation, each IRAC image is corrected by dividing it by a map of these variations, called a flat field.

Observations are taken of approximately 20 pre-selected regions of high-zodiacal background chosen to be as free as possible of stars and Galactic cirrus. Between 80 and 100 dithered frames of 100 seconds in each channel are taken. These observations are processed in the same manner as science data and then averaged with outlier rejection. Additionally, objects are detected during the averaging and rejected. The flat fields are then normalized to one. The flat field calibration thread produces a library of the flat fields throughout each campaign since a flat field is taken at the beginning and end of an observing campaign, and every time darks are taken. Analysis of the first two years of flat field data taken during nominal operations indicates that the flat field response of IRAC is fixed to within our ability to measure. As a result, "super skyflats" were made from the first two years of data, and were loaded as static calibration data. Thus, as of pipeline version S13, all IRAC data are flattened with the same "two-year skyflats." See Section 4.2.2 for further discussion of the flat fields.

3.2.5 Cosmic Ray Detection

Individual frames are analyzed for probable radiation hits (cosmic rays), and the results appear as a flag in a radhit image. This is computed by a median filtering technique. Input images are read in, and a median filter is applied. The difference between the input image and the median-filtered image is then computed. Pixels above a specified threshold (i.e., are "pointier" than is possible for a true point source) are then flagged in a mask image (bit 9 of dmask is set when a pixel is suspected to be hit by a cosmic ray). The online detection uses a very conservative (high) threshold so that real point sources are not rejected. See the discussion in Section 4.5 for more information on outlier rejection.

3.2.6 Flux Calibration

IRAC flux calibration is tied to a system of celestial standards. These standards are measured at regular intervals. For details of the photometric calibration see §5.1.1.

IRAC data are calibrated in units of MJy/sr. This is accomplished by multiplying the image by a conversion factor. The flux conversion factor is in the image header keyword FLUXCONV and mirrors the values given in Reach et al. (2005).

3.3 IRAC Post-BCD Pipeline

3.3.1 Pointing Refinement

To improve the $\approx 0.''5$ blind pointing, a pointing refinement is run in which the point sources are identified in the IRAC frames and correlated with 2MASS source positions. The pointing refinement typically improves the positional error to $< 0.''3$ and removes any systematic offsets.

First, point sources are extracted from the pipeline-processed mosaics and transformed to RA and DEC using the transformations derived from the current pointing. If there are less than five sources in an image, then there will be no refinement for that BCD.

A comparison is made of the position and flux of each point-source match found in the 2MASS point-source catalog. The new translational and rotational reference frame can be computed from the differences and uncertainties, and a refinement is made of the celestial pointings and angles of each BCD in the AOR used for the mosaic. These refined values are written to the end of the FITS headers as RA_RFND, DEC_RFND, including many others with _RFND as an indicator of "refined" pointing.

3.3.2 Superboresight Pointing Refinement

Pointing refinement operates on each IRAC channel independently. This often results in poor pointing solutions for channels 3 and 4, in which stellar fluxes are lower and the background higher than in channels 1 and 2. We have therefore developed a technique which combines the results of pointing refinement in channels 1 and 2 and applies it to all four channels using the known offsets between the IRAC fields of view. This improved pointing solution is derived during campaign reprocessing. The results of the pointing refinement from the first run of the post-BCD pipelines are averaged, and the correction derived from this is applied to the boresight pointing history file (which contains the pointing estimate derived from the spacecraft telemetry and which provides our initial pointing estimate). This corrected pointing history file (the "superboresight" file) is then applied to the BCD at the pointing transfer stage of the BCD pipeline. The superboresight RA and DEC estimates are recorded in the CRVAL1 and CRVAL2 FITS keywords, and the position angle estimate is recorded in the CD matrix keywords. The uncorrected RA and DEC are retained, but called ORIG_RA, ORIG_DEC, as

is the pointing refinement solution for each frame (as RARFND and DECRFND). Note that to use the superboresight solution, USE_REFINED_POINTING = 0 should be set in the MOPEX namelists.

Superboresight is still being fully tested, and will be implemented as a patch to the S13 software build, but it is likely that most data processed or reprocessed in S13 will use it. At the time of writing it is unclear though whether all data processed under S13 will have the supeboresight applied. Users should check for the presence of the ORIG_RA, ORIG_DEC keywords to see if it has been applied to their data.

Currently the implementation applies the superboresight correction to the short and medium frames of HDR data, but a small improvement can be made by copying the CRVAL1, CRVAL2 of the long frame into the short and medium frame headers, as the pointing refinement solutions for the long frames on which the supeboresight solution is based are usually more accurate. This will be automated in a future software release.

3.3.3   Mosaicking and Coaddition

Using the refined coordinates, individual IRAC BCDs from a given AOR are reconstructed onto a larger field (mosaicking), and overlapping frames are averaged together to achieve a higher S/N. Outlier rejection is performed on sets of overlapping pixels. Because Spitzer observations cover such a large area, individual BCDs are remapped onto a common grid with a technique similar to 'drizzle' (Fruchter & Hook 2002). The pixel size in the mosaics produced by the pipeline is exactly 1.2 arcseconds × 1.2 arcseconds.

The masks are used in the coaddition in such a way that the pixels previously flagged as bad (for example, hot or dead pixels) are rejected before the averaging process. Cosmic rays are rejected at this point via the outlier rejection algorithm.

Users will receive a single coadded image per channel, and per AOR. It will be accompanied by a coverage map and an uncertainty file per channel. For more details about post-BCD data and the processing of IRAC images, see Section 7.

Chapter 4

Detector Artifacts and Cosmic Rays

4.1 Introduction

Observers receiving data that appear flawed or scientifically invalid should attempt to identify the cause and severity of these problems, and then send a note to the *Spitzer* Helpdesk (help@spitzer.caltech.edu). Notifying the SSC should be done within two months of the retrieval of the proprietary data. Observers may request that the pipeline be re-run on their data if the raw data are OK but it appears that the pipeline has failed. If the raw data are unusable, the problem is probably unrecoverable. In that case, observers can request an observation be deemed "failed" and be repeated by sending a note to the *Spitzer* Helpdesk. Some of the common artifacts in IRAC data are discussed in this chapter. Most of these can be mitigated by a judicious quality inspection of the data, and/or further processing of the BCDs. Note that many of these artifacts are quite commonly seen in IRAC images. Unless they clearly compromise the scientific use of your data, they do not justify the reprocessing or reobservation request.

The most common artifacts are as follows. Stray light from point sources should be masked by hand and is not normally grounds for a reobservation request. Persistent images from a bright source observed as part of the observation are similarly not normally grounds for a reobservation request. In some cases, however, persistent images from a preceding observation may be found. One way to check this is by inspecting a median of all images in an AOR. Another possible flaw in the observations would be an exceptionally high radiation dosage. The nominal rate is 1.5 hits per array per second, and the radiation hits range from single pixels to connected streams (and occasionally small clouds of secondaries). High particle hit rates occurred following one solar flare during the In-Orbit Checkout, and one in Nominal Operations. In the latter event, several hours of science data were rendered useless because of the large number of hits in the images. Objects that are bright enough leave muxbleed trails and can generate pinstripe patterns over large parts of the image, and offsets along the columns and rows containing the bright source. Ghosts from internal reflections within the filters can be seen in almost every channel 1 or 2 BCD, and more ghosts in all channels are noticeable from bright objects.

We begin with a discussion of the basic characteristics of the dark frames and flat fields that affect every image. We follow with a discussion of electronic artifacts. These effects arise from the inherent nonlinearity of the detector diodes and saturation of either the detector well, transistors in the mux, or the analog-to-digital converter (ADC) in the warm electronics; crosstalk within the mux or warm electronics; or from inductive coupling to currents in spacecraft cables. Most electronic effects have a short persistence, but image persistence, which is also nonlinear in photon fluence, can last seconds, minutes, hours, or even weeks. Next we have a section on optical artifacts, which include stray light or ghosts from sources within or outside the FOV. Finally, we discuss the effects of cosmic rays and solar protons on IRAC observations.

4.2 Darks, Flats, and Bad Pixels

The true median dark currents, due to nonzero leakage resistance or recombination in reverse-biased detector diodes, are very small compared to the current from the background at the darkest part of the celestial sphere. Labdarks, which were measured with the cold IRAC shutter closed, with zero photon flux, are not zero, and have significant pixel-dependent offsets, usually positive, that depend on the frame time and the Fowler number, as well as the history of readouts and array idling over the previous several hours. Channel 3 is by far the most extreme case, in which, for example, a 100 second (Fowler-16) frame can be offset as much as 370 DN (median), or the equivalent of 1400 electrons at the integrating node, with no light incident on the array. The signal from the darkest background in a 100 second frame in channel 3 is only about 1000 electrons. Channel 1 has much smaller offsets, but the sky is so dark that the offsets are often larger than the background signal. Channel 2 has very small offsets, which are less than the background signal except in short frames or certain frames immediately following a change in integration time. The background in channel 4 is so large that the offsets are almost negligible except in very short integrations. There is no measureable excess noise from the offset itself: the noise is not the square root of the equivalent number of charge quanta on the integrating node. This is because the offset arises from the redistribution of charge within the mux in which the associated currents and capacitance are much greater than in the detector diodes. However, imperfect correction of the first-frame effect does increase the uncertainties in BCD frames. The uncertainty scales with the size of the offset and its small-scale spatial nonuniformity. Only in channel 3 does it significantly increase the total pixel noise.

We can break down the offset into contributions beginning with the largest spatial scale down to the smallest. In this view, the largest part of the offset is uniform over the array, followed by the contribution of a few spatial gradients, and some pinstriping that repeats every 4 columns (due to the 4 array outputs), with a few columns with odd offsets (due to hot pixels or parts of the mux), and weakest of all, pixel-to-pixel dependent offsets.

There are some very obvious features imposed on the true offset, due to a relatively small number of hot pixels, and mux glow. Hot pixels usually appear bright, and in such cases one can see a trail of muxbleed (in raw images) or a pinstripe pattern in InSb arrays, or the bandwidth effect (in Si:As arrays) following the hot pixel. These pixels have high dark currents and are usually isolated but sometimes in a clump. "Dead" pixels are really just very hot pixels, so hot that they saturate before the first pedestal sample. In a BCD image, hot pixels do not appear bright because they have been canceled by the labdark or skydark subtraction. Most hot pixels appeared after launch and are the result of hits by energetic nuclei. By annealing the arrays, we restore most pixels that get activated. Some of them cannot be restored, and thus they become "permanent" hot pixels. Some pixels jump randomly from normal to high dark current and back, dwelling in one state for anywhere from a few minutes to weeks, so they may not be canceled by a skydark subtraction. These are IRAC's "rogue pixels." The IRAC "static" bad pixel masks are updated when significant changes in the permanent bad and/or hot pixels occur.

Areas of mux glow are visible in the labdark and images. Electrons and holes recombine in diodes in the mux, allowing current to flow. Photons emitted in the recombination are detected in the InSb or Si:As detector above or near the source of the glow. Most prominent is the glow from the four output FETs visible only in Channels 3 and 4 (the Si:As arrays). These are semicircular areas about 17 pixels in radius located near column 256, row 30 at the right edge of the image. The glow is most obvious in long frames. Another glow region is visible along the last few rows in all 4 channels; it comes from the unit cell FETs. Currents flow through all the unit cell FETs in the last row which is left selected during the integration, so the glow is particularly bright in the last row itself. The 3rd and faintest glow region is along the left edge (column 1) of channel 3. Detected glows have shot noise, which can exceed the background

Table 4.1: Definition of bits in the 'pmask'.

Bit	Condition
0	Not set
1	Not set
2	Not set
3	Not set
4	Not set
5	Not set
6	Not set
7	Dark current highly variable
8	Response to light highly variable
9	Pixel response to light is too high (unacceptably fast saturation)
10	Pixel dark current is too excessive
11	Not set
12	Not set
13	Not set
14	Pixel response to light is too low (pixel is dead)
15	[reserved: sign bit]

noise along the last row and in the brightest parts of the semicircular areas. Pixels are masked in these areas where the noise significantly degrades sensitivity in 100-second frames.

4.2.1 Bad Pixels

Included with the BCD data are three masks, a semi-static mask (the "pmask") which containing permanent or semi-permanent bad pixels and regions, and which is the same for all BCDs in a given AOR and channel, the "dmask" which contains bad pixels specific to any one BCD, and the "imask" which is a more user-friendly version of the dmask. Some of the information contained in the dmask is not relevant to the final data quality and can be ignored. All of the bits set in the imask indicate pixels that have been compromised in some fashion. Not all of the imask bits are set by the pipeline, but some bits are placeholders for post-BCD processing of data artifacts. The higher the order of bit set in the imask, the more severe the effect on data quality. Mask values are set as powers of two, and summed together for each pixel. Any pixel with a bit set in the pmask is suspect. Bits 5 and 8–14 in the dmask are indicators of poor quality data. The remaining dmask bits are informational but do not indicate that a pixel is bad. There are currently five sets of pmasks to account for variations in bad pixels with time. Tables 4.1, 4.2, and 4.3 show what each bit value corresponds to.

4.2.2 Flat Field

Individual pixel-to-pixel gain variations are corrected by means of a pixel-to-pixel gain map commonly known as a "flat field." IRAC flats are derived by making highly dithered observations of one of approximately 20 fixed locations in the ecliptic plane, specifically chosen to be as free of stars and extended cirrus emission as possible, and where the zodiacal light provides a uniform illumination. The data are processed much like science data and then averaged with outlier rejection. Additionally, since stars, asteroids and galaxies are a significant contaminant in the data, an object detector is used to find and then explicitly reject them during the averaging. The flats are normalized to a median of one. New flat field measurements are made every time the instrument is turned on.

Table 4.2: Definition of bits in the "dmask."

Bit	Condition (pipeline module bit is set by)
0	Muxbleed correction partial or not applied (muxbleedcorr)
1	First/Inter-frame correction not applied (ffcorr)
2	Pixel contaminated by optical ghost (not yet implemented)
3	Pixel contaminated by electrical ghost (not yet implemented)
4	Sign wraparound correction was applied (iracwrapcorr)
5	Latent-image flag (latimflg)
6	EBWC could not be applied (iracewbc)
7	Flat field applied using questionable value (flatap)
8	Flat field could not be applied (flatap)
9	Cosmic ray detection (detect_radhit)
10	Data saturated (fowlinearize); HDR data saturated (predictsat)
11	Data known to be bad before dmask generation (cvti2r4)
12	Linearization not applied (fowlinearize)
13	Pixel contaminated by electrical crosstalk (banding flag; not yet implemented)
14	Data missing in downlink (cvti2r4)
15	[reserved: sign bit]

Table 4.3: Definition of bits in the "imask."

Bit	Definition
0	reserved for boolean mask (or if best practice bits set, data quality)
1	reserved for future use
2	optical ghost flag (to be set by post-BCD tool)
3	stray light flag (to be set by post-BCD tool)
4	donut detection flag (to be set by post-BCD tool)
5	muxbleed flag (to be set by post-BCD tool)
6	banding flag (to be set by post-BCD tool)
7	pulldown flag (to be set by post-BCD tool)
8	crosstalk flag (bit 13 of dmask)
9	radhit (bit 9 of dmask, single frame radhit detection)
10	latent flag (bit 5 of dmask)
11	not flat-field corrected (bit 8 of dmask)
12	data not very linear (to be set by future pipeline module)
13	saturated (bit 10 of dmask)
14	data bad and/or missing (bit 14 OR bit 11 of dmask)

Analysis of data from the first two years of operations has showed that the flat field response of IRAC is unchanging at the limit of our ability to measure. As a result, so-called "super skyflats" were generated from the first two years of data. The superskyflats are shown in Figure 4.1.

These flats are extremely low-noise, with stochastic pixel-to-pixel uncertainties of 0.4%, 0.2%, 0.2%, and 0.05% in channels 1 through 4, respectively. This is smaller in amplitude than the intrinsic pixel-to-pixel scatter in the gain. Furthermore, because the super skyflats are derived from data over many parts of the sky, with many dithers and rotations of the telescope, they are substantially free of errors arising from gradients in the zodiacal background, or from residual contamination by stars and galaxies. Currently all IRAC data is reduced with the same set of super skyflats.

Large-scale gradients corrected by the flats are on the order of 10%-15%. Systematic errors in the flats are due to the gradient in the zodiacal background, and straylight removal errors. The former is expected to be very small based on results from other missions (Abraham et al. 1997, ISOPHOT 25 μm). Diffuse stray light is a significant contaminant in the raw images at the $\sim 5\% - 10\%$ level. This diffuse light looks like a 'butterfly' across the top of the InSb detectors in channels 1 & 2, or a 'tic-tac-toe' pattern in channels 3 & 4. It is always present, resulting from scattering of the zodiacal background onto the detectors. In both the skyflats and the science data, a model of the straylight has been subtracted, but this leaves a residual pattern on the order of 1% which contaminates the flats. These errors are substantially ameliorated by dithering – errors will decrease as $\sqrt{N}$, where N is the number of dithers, and will quickly become very small relative to other uncertainties.
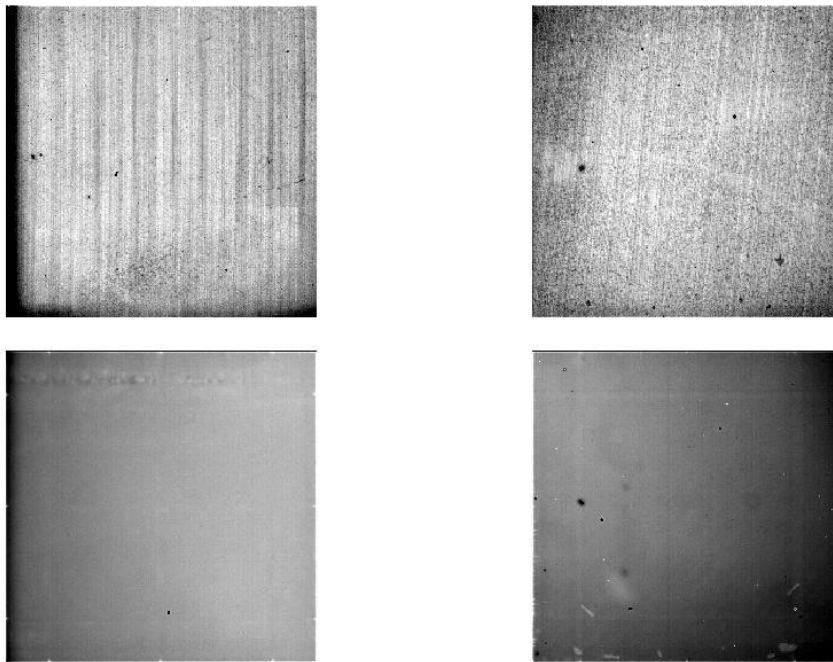
Figure 4.1: Superskyflats for IRAC. These were made by combining the flat fields from the first two years of operations. Note the dark spots on the 4.5 and 8.0 μm flats. These are due to impurities on the pick-off mirror. These pixels are about 30% dimmer than the surrounding flat pixels. Flat fielding in the pipeline fully corrects for these dark spots in the data.

There is a slight level of vignetting of the IRAC arrays, most evident in Channels 1 and 3. At no point are the arrays fully blocked. However, their responsivity falls to roughly 50% at the

left-most edge. The vignetting only extends for 10–15 pixels. The vignetting is compensated for by the flat-fielding, and results primarily in an increase by at most $\sqrt{2}$ of the noise in the affected pixels.

Finally, one should note that the flatfields are generated from a diffuse, extremely red emission source. While the resulting flats perfectly flatten the zodiacal background, they are *not accurate* for compact objects with different spectral slopes, the most obvious examples being stars. Please see the section on array-location-dependent corrections (Section 5.3).

4.3 Electronic Artifacts

4.3.1 Saturation & Nonlinearity

There are four places in the electronics where a pixel may saturate: the detector diode, the unit cell source-follower in the Read-Out Integrated Circuit (ROIC), the output source-follower in the ROIC, and the analog-to-digital converter (ADC) in the warm electronics. In most cases, it is the ADC that saturates first, at 0 or 65,535 units. ADC saturation produces a discontinuity in the second derivative of the measured Fowler DN versus the flux. The other saturations are smooth, with no discontinuity. In the other cases, depending on the channel, the detector diode may saturate before or after the source-followers.

For observations of very bright sources where the pedestal reads are saturated, the DN for a particular pixel and frame time and Fowler number is double-valued for fluxes between 0 and some very high value. Thus, for a bright, strongly saturated point source, the DN will increase from some low number away from the source to some maximum value of 35,000 to 47,000 DN, and then decrease to a small, usually negative number, at the center. The image looks like a bright doughnut with a dark center. In principle, for any source for which we already know the spatial variation of its intrinsic surface brightness, we can determine whether the pixel is above or below saturation, and therefore, its flux. In practice, we do not know the gains of the source-followers very well near saturation, nor do we know enough about the detector diode saturation, to make a good estimate of the flux. Therefore, we flag pixels which are above the range of our linearization correction. For point sources, a rough estimate of the flux in the saturated pixels can be made by fitting the wings of the PSF to the linearized pixels in the BCD image. Pixels which saturate in the pedestal reads are not flagged by the SSC pipeline and need to be identified. The short frame time data in HDR observations can be used to recover saturation in the long frame time data. Pixels with significant flux in the short frame that are above the saturation limit in the long frames can be substituted. An IDL routine to perform this task is available at the SSC contributed software site. This replacement is accurate to about 10% at the peak of bright sources as the ~ 0.2 arcsecond jitter of the telescope coupled with pixel phasing in channels 1 and 2 and charge diffusion in all channels will cause the measured flux densities between short and long frames to vary.

The IRAC arrays are slightly nonlinear at all signal levels. At levels above 30,000 DN (in the raw data) the response is low by several percent. As part of pipeline processing, the data are linearized based on ground calibrations (which have been verified in flight) of this effect. The BCD data are linear to better than 1% up to about 90% of full well, which is defined to be the level where we no longer can fully linearize the data, and at which saturation, by definition, begins. Below 20% of full well the nonlinearity in the raw data is negligible.

4.3.2 Muxbleed (InSb)

Multiplexer bleed, or "muxbleed," appears in IRAC channels 1 and 2 (3.6 and 4.5 μm). It looks like a decaying trail of pixels, repeating every 4th column, with enhanced output level trailing a

bright spot on the same array row. The effect can wrap around to subsequent rows, but it does not wrap from the last row to the first. Since columns are read simultaneously in groups of 4, one for each mux output, the next pixel read out on any single output is 4 pixels to the right, in array coordinates. As the BCDs for channels 1 & 2 are flipped in the y-direction when compared to the raw images, the read direction is top to bottom for BCDs and muxbleed-triggering pixels will affect rows beneath the source. Muxbleed is usually accompanied by a pinstripe pattern (every 4th column) that may extend over part of the image preceeding or following the pixel. It is caused by a slow relaxation of the mux following the momentary disequilibrium induced when a bright pixel's voltage is placed on an output FET during pedestal and signal reads. Although the pixel rise and fall times are fast (2.6 and 1.0 μsec, respectively) compared to the 10 μsec time to clock the next pixel onto an output, longer relaxation times are involved for an output FET to fully recover after the voltage from a bright pixel is briefly impressed on its gate. The decaying trail has a time constant of tens of μsec, and the pinstripe, tens of seconds. In BCDs produced by pipeline versions prior to S13, the pinstripe patern from muxbleed was complicated by a de-striping step in the pipeline in the darkdrift module. This often caused pinstriping to appear over an entire image. Beginning with pipeline version S13, we turned off the de-pinstriping in channels 1, 2, and 4, but left it on for channel 3.

Stars, hot pixels, and particle hits can generate muxbleed, and the characteristics of the pinstripe depend on frame time and Fowler number. Hot pixels may show muxbleed in a raw image, but in the BCD the muxbleed induced by hot pixels may not be present because it was canceled in either the labdark subtraction or in the skydark subtraction. The pinstripe pattern is nearly constant in areas of a single image that do not contain a saturating star, particle hit, or hot pixel. The characteristics of muxbleed from particle hits depend on when the hit occurs within the frame.

Muxbleed was characterized long before the launch of *Spitzer*, but is not understood well enough to be fully corrected in the pipeline. The correction as it is performed in the present pipeline is described in §3.2.2. The pinstripe is strongest in channel 2, particularly in 12 second frames. In a channel 2 mosaic, even with overlap correction, there may appear to be bright and dark patches everywhere, about the size of one frame or part of a frame. Upon close inspection, though, individual patches are revealed as areas of nearly constant pinstripe pattern that runs between edges of the array, bright stars, hot pixels, and particle hits.

Although the amplitude of the effect decays as one moves away from the bright spot, this decrease can not be fully described by a simple function. In general, the muxbleed decays rapidly within 5–10 reads and plateaus at a roughly constant value. The functional form of the muxbleed is frame time dependent and has not been characterized in detail. In particular, the amplitude does not scale linearly with the flux at the brightest pixel or the integrated flux of the triggering source. In addition, since in many cases the pixel corresponding to the bright star is saturated, predicting the magnitude of the muxbleed effect based on the brightness of the offending star is difficult. There is no threshold below which muxbleed does not occur. Usually it is visible in a single frame only when a pixel is near or above saturation. However, as more frames from the same position on the sky are averaged together, the muxbleed from fainter stars rises above the noise.

Some attempts have been made to fit empirical profiles of muxbleed and correct it to below the rms noise level of the image. This can be done by fitting a decaying function and scaling it either by the brightness of the offending pixel (below saturation) or that of certain pixels some distance away (if the offending pixel is beyond saturation), or by fitting a model of the muxbleed to the affected pixels using an estimate of the true sky in the affected pixels. In the latter case, a model can be derived from the user's data to tune the solution to a particular problem. The scale factor must be fitted to each point source above 3/4 full well. Both algorithms are somewhat successful and are available as contributed software at the SSC website. An example of the
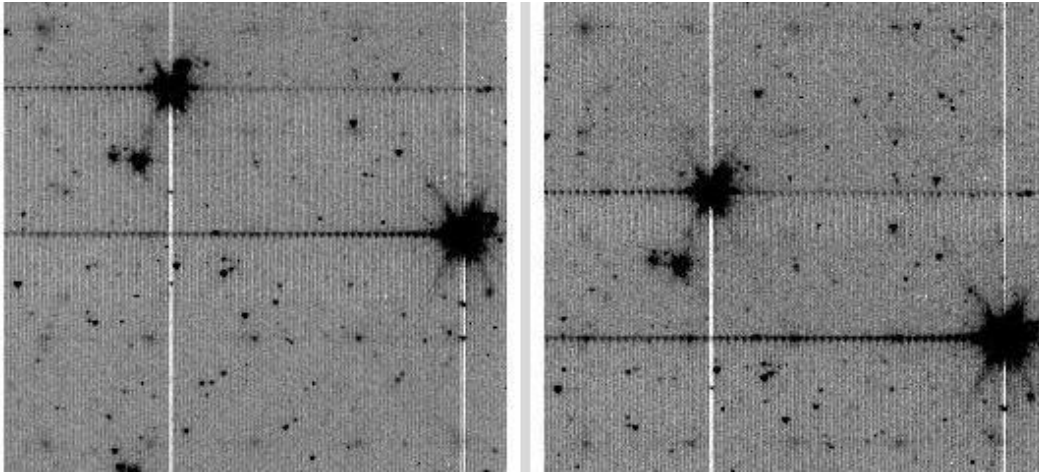
Figure 4.2: Images showing muxbleed effect (the horizontal line on both sides of a bright stellar image). The pixels on the left side of the bright source are pixels on rows following the row where the bright source was located (and have wrapped around in the readout order of the array). The vertical (white) lines are due to the so-called "column pull-down" effect, discussed in Section 4.3.4. These are 12-second BCD frames in IRAC channel 1, taken from IRAC program pid = 618, AORKEY = 6880000.

result of such processing to correct for muxbleed is shown in Figure 4.3. It can be seen that at least cosmetically the effect can be greatly reduced without introducing new artifacts. However, depending on the nature of the bright source and its level of saturation, the results are not uniform, and even in the best cases some residuals are still visible.
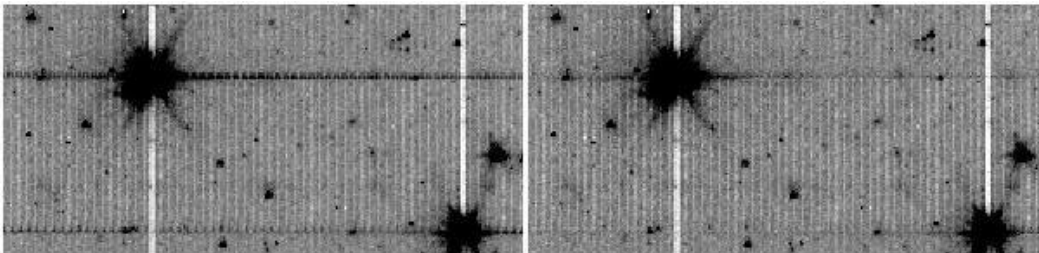

Figure 4.3: Experimental removal of the muxbleed effect. The image on the left is before and the one on the right is after the correction. These are in-flight channel 1 data, taken from IRAC program pid = 618, AORKEY = 6880000.

Currently there is no code available to correct for pinstriping that is associated with muxbleed (see Figure 4.3). A good first step would be to match the median values of each output in the affected region.

4.3.3 Bandwidth Effect (Si:As)

The bandwidth effect appears in IRAC channels 3 and 4 (5.8 and 8.0 um). It looks like a decaying trail of pixels 4, 8, and 12 columns to the right of a bright or saturated spot. Only in the most highly saturated cases is the effect visible 12 columns to the right. A typical case for a star is shown in Figure 4.4. The effect is due to the fact that inside the ROIC the maximum voltage slew rate is limited, so charge on the output bus can not be drained fast enough for the output to settle to the value for a dark pixel that follows a bright pixel, or vice versa, in the 10 μsec or 20 μsec at which times the next two pixels (4 and 8 columns to the right) are read

out. A smaller, additional delay comes from charging or discharging the cables from the array to the warm electronics. The effect is nonlinear except in the weakest cases. The output FETS in the Si:As arrays do not have the long recovery time that causes the long muxbleed trails and pinstriping in the InSb arrays, in part because the voltage swings have the opposite sign. The bandwidth effect presumably affects the first two or three pixels read out after a bright pixel in the InSb arrays as well, but for InSb, we have included the bandwidth effect as part of the overall "muxbleed" effect. It is much better behaved in InSb because the voltage swings are smaller and the slew rates are faster. A rare case which gives rise to a bizarre image is shown in Figure 4.5. Here, an extremely saturated star saturates an area in the last 4 columns. The bandwidth effect appears in the first 12 columns, making it appear as if the right edge of the image was cut and pasted onto the left side of the image. Because of the details of the array clocking, part of the unsettled signal appears in both the same row and the next row.
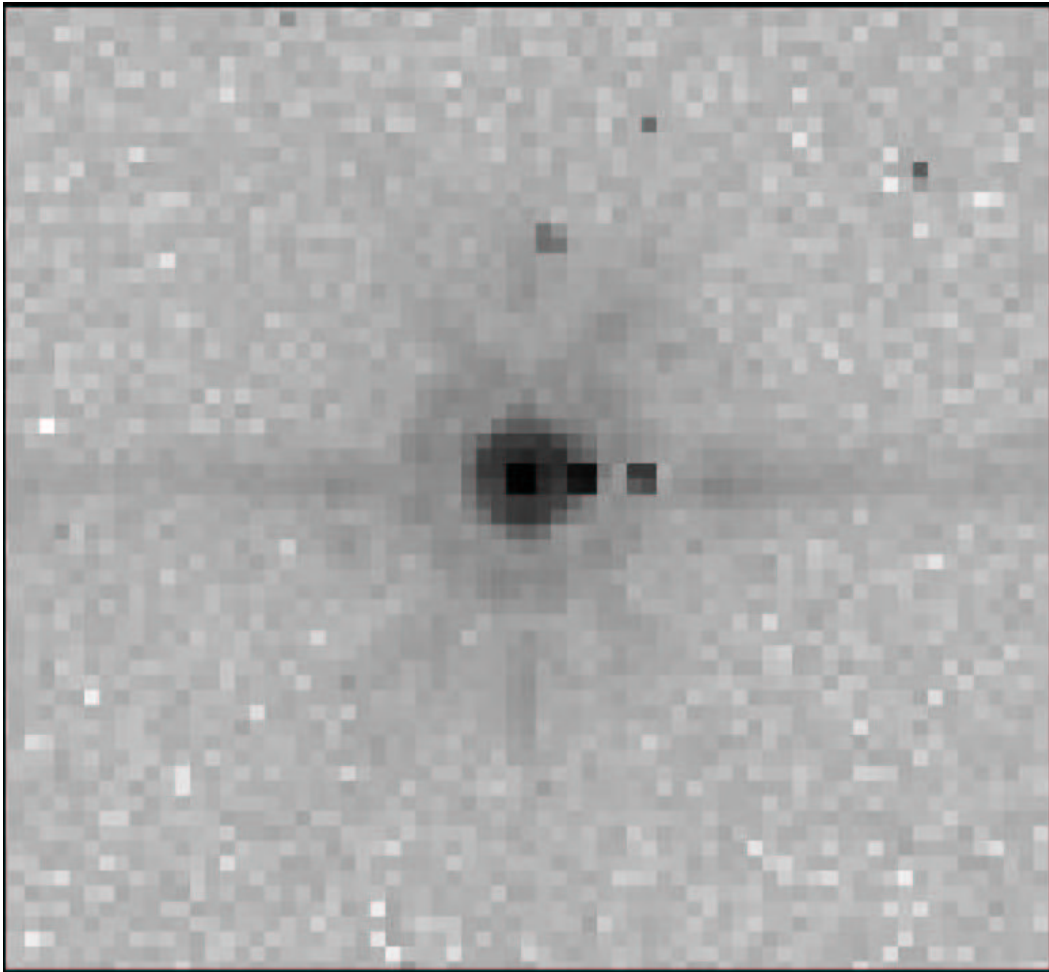
Figure 4.4: A typical bandwidth effect trail in channel 4, in a 30 second frame. These data were taken from program pid=1154, AORKEY = 13078016.

4.3.4 Column Pull-up/Pull-down

When a bright star or cosmic ray on the array reaches a level of approximately 35,000 DN, there is a change in the intensity of the column in which the signal is found. In channels 1 and 2, the intensity is reduced throughout the column (thus the term "column pull-down"); see Figure 4.6. When the effect occurs, it shifts the intensities of the pixels above and below the

Figure 4.5: The bandwidth effect when a bright object is in the last 4 columns. IRC+10216, strongly saturated, is just off the right side of the channel 3 array. Even the filter ghost is saturated. The bandwidth effect appears on the left side of the array. These data were taken from program pid = 124, AORKEY = 5033216.

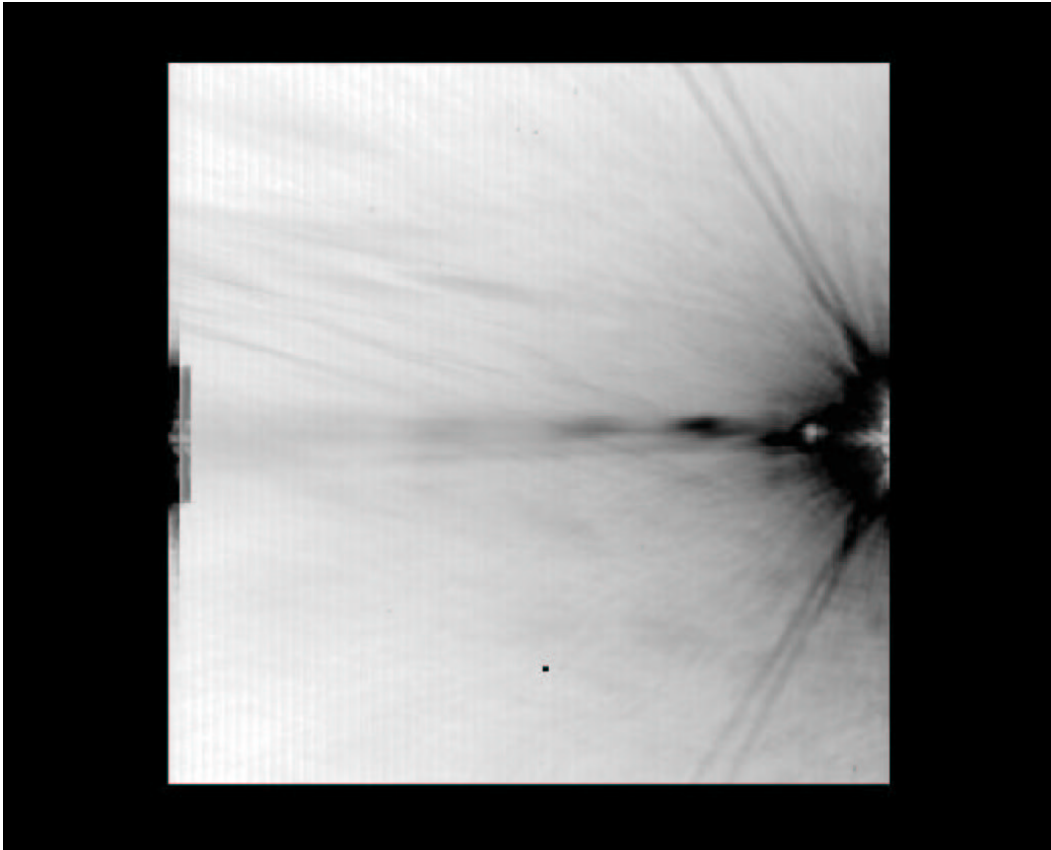position of the "guilty" source, within the same column. This effect is limited to the brightest sources. The amplitude of the column pull-down does not scale linearly with the flux of the source or the brightest pixel. The effect appears to be constant on either side of the source and algorithms which fit separate DC offsets above and below the source should be effective. Cosmetic corrections are partially successful. One, provided by the GOODS Legacy team, takes the median of each column, identifies columns that deviate from the local average by more than some threshold, then adds back in a constant to the apparently affected columns. The code does not currently work in fields with extended emission. A more general algorithm which estimates the "true" sky value for affected pixels and fits DC offsets is also available for observations of more structured emission. This algorithm is posted on the SSC website as contributed software (see also the link to the pull-down corrector IDL code in the SSC Post-BCD web pages).

4.3.5 Row Pull-Up

In addition to muxbleed in channels 1 and 2, there may be electronic banding, which is manifest as a positive offset for rows that contain bright pixels. This effect is at least an order of magnitude smaller than muxbleed, and is currently under investigation. Electronic banding is more significant in channels 3 and 4 but it is not as significant as the optical banding in those channels (see Section 4.4.2). The IRAC artifact mitigation software at the contributed software site can mitigate these effects. The software finds instances of pull-up and banding and fits the
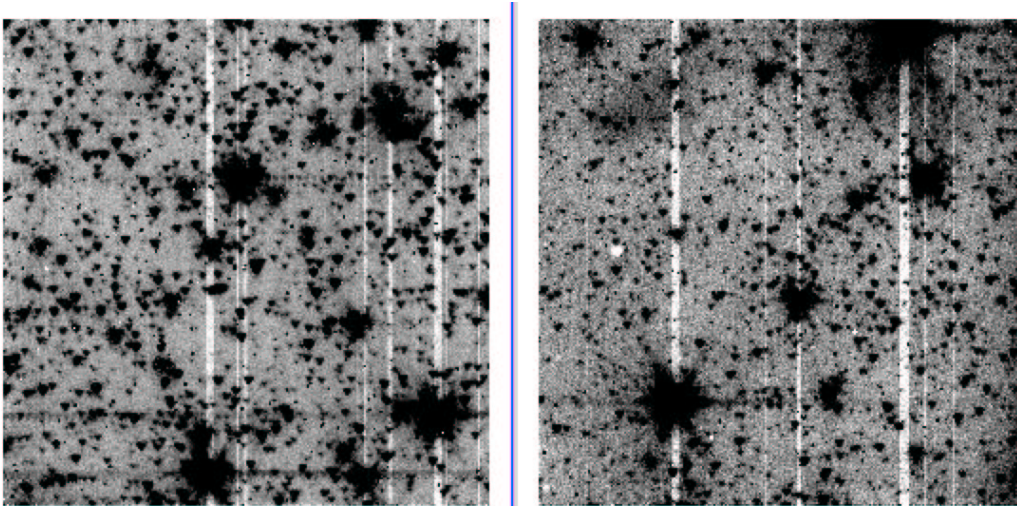
Figure 4.6: IRAC channel 1 (left) and channel 2 (right) observations of a crowded field with column pull-down apparent from the brightest sources. Note that the brighter sources affect a larger number of columns. These data were taken from program pid = 613, AORKEY = 6801408.

DC offsets on either side of the triggering source to them.

4.3.6   Full-Array Pull-Up

In all four arrays, there is also an effect where the entire array can be uniformly offset in the presence of a bright source. This effect probably occurs also for uniform illumination, but it is not obvious. The effect, seen in channel 3 and 4 IBC arrays, is sometimes called "droop." Array pull-up can be mitigated by matching the background level of the affected BCD with that of the overlapping, unaffected BCDs.

4.3.7   Inter-Channel Crosstalk

We have detected electronic crosstalk between channels only in the brightest sources that have been observed. All four channels are read out simultaneously, except for the 100 second frames in channel 4, for which two 50-second frames are taken instead of the long integrations in the other channels, because IRAC is background limited in channel 4. When a source falls on a pixel of one array, crosstalk may occur in the same pixel location in the other arrays, or in the next pixel read out. The crosstalk appears as a combination of either a positive or a negative offset in the same pixel, and the derivative of the signal in the same or previous pixel. As the source is dithered, the crosstalk follows it, and therefore crosstalk appears in the mosaics. It is so weak that we have detected it so far only in channel 3, when the source is in channels 2 and 4, and in channel 4, when the source is in channels 1 and 3. Figure 4.7 shows parts of the mosaics from the off-beams in a dithered observation of a very bright star. The star was observed in channel 3.6/5.8 FOV first, so there are residual images in channels 1 and 3 from the bright star. The residual images appear as a diffuse glow near the center; they are a combination of the residual images of a very strongly saturated star observed with a Reuleaux dither pattern, thus effectively smoothed by outlier rejection. The crosstalk appears in channels 3 and 4 as a partial dark ring with a bright core.
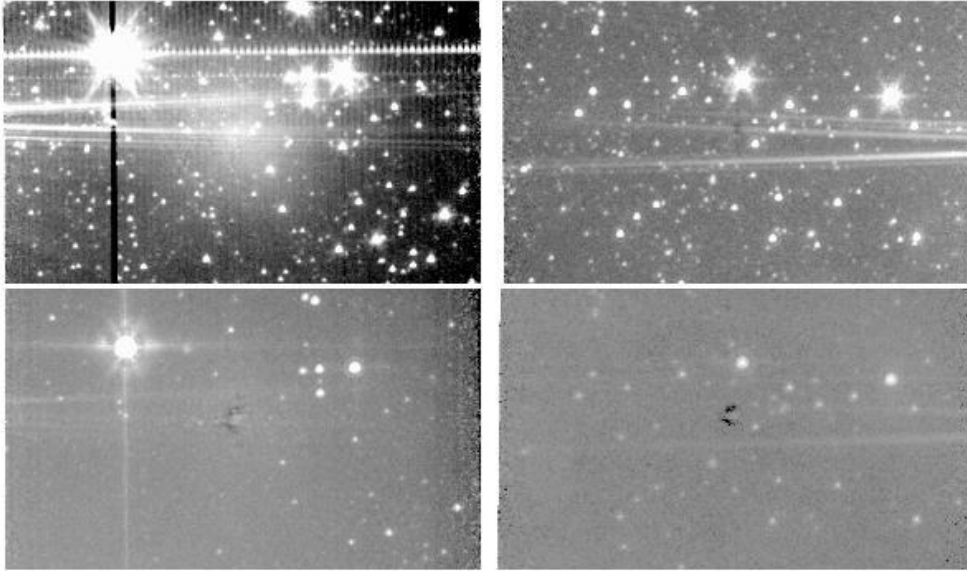
Figure 4.7: Channels 1 and 2 (top) and 3 and 4 (bottom) showing inter-channel crosstalk (dark spots near center of lower panels).

4.3.8 Persistent Images

The terms "persistent image", "residual image", and "latent image" are used interchangeably to describe the contamination of an IRAC image by a bright source from a previous exposure. Tests performed during In-Orbit Checkout (IOC) revealed that there are both short-term persistent images, with time scales of order minutes and which are present in all 4 arrays, and longer-term persistent images in channels 1 and 4. The short-term persistent images were known before launch, and extensive calibrations and data analysis were performed to characterize them. The pipeline produces a mask (bit 5 of the Dmask) for each image that indicates whether a bright source seen by a previous exposure would have left a persistent image above three times the predicted noise in the present frame.

The longer-term persistent images were discovered in flight. In channel 1, the persistent images are generated by stars as faint as $K = 13$ (in very long stares). They can be generated by any long dwell time with a bright star on the array, whether or not the array is being read out. They were first noticed during a high-gain antenna downlink, when IRAC was left at a fixed position viewing the Galactic plane (by chance) for 45 minute. The persistent images do not have the same size as a direct point source; they are significantly more diffuse (looking more like the logarithm of the point-spread function). The channel 1 long-term persistent images have time scales of order 6 hours, and they decay gradually. The cause of these persistent images has been identified as a known feature of the flight array (broken clamp) that cannot be fixed.

The longer-term persistent images in channel 4 are induced by bright mid-infrared sources or bright stars. The channel 4 persistent images have very unusual properties: they have lasted for as long as two weeks, they can survive instrument power cycles, and they do not decay gradually.

The SSC has instituted a pro-active and highly successful method of eliminating persistent images. Channels 1 and 4 are temporarily heated, or "annealed," briefly, with a small current running through the detector. The arrays are annealed after every telemetry downlink, which erases any persistent images built up during the downlink or during the previous 12-hour period of autonomous operations (PAO). This strategy, combined with scheduling known bright object observations immediately before downlinks, greatly decreases the possibility that preceding

observations will produce persistent images.
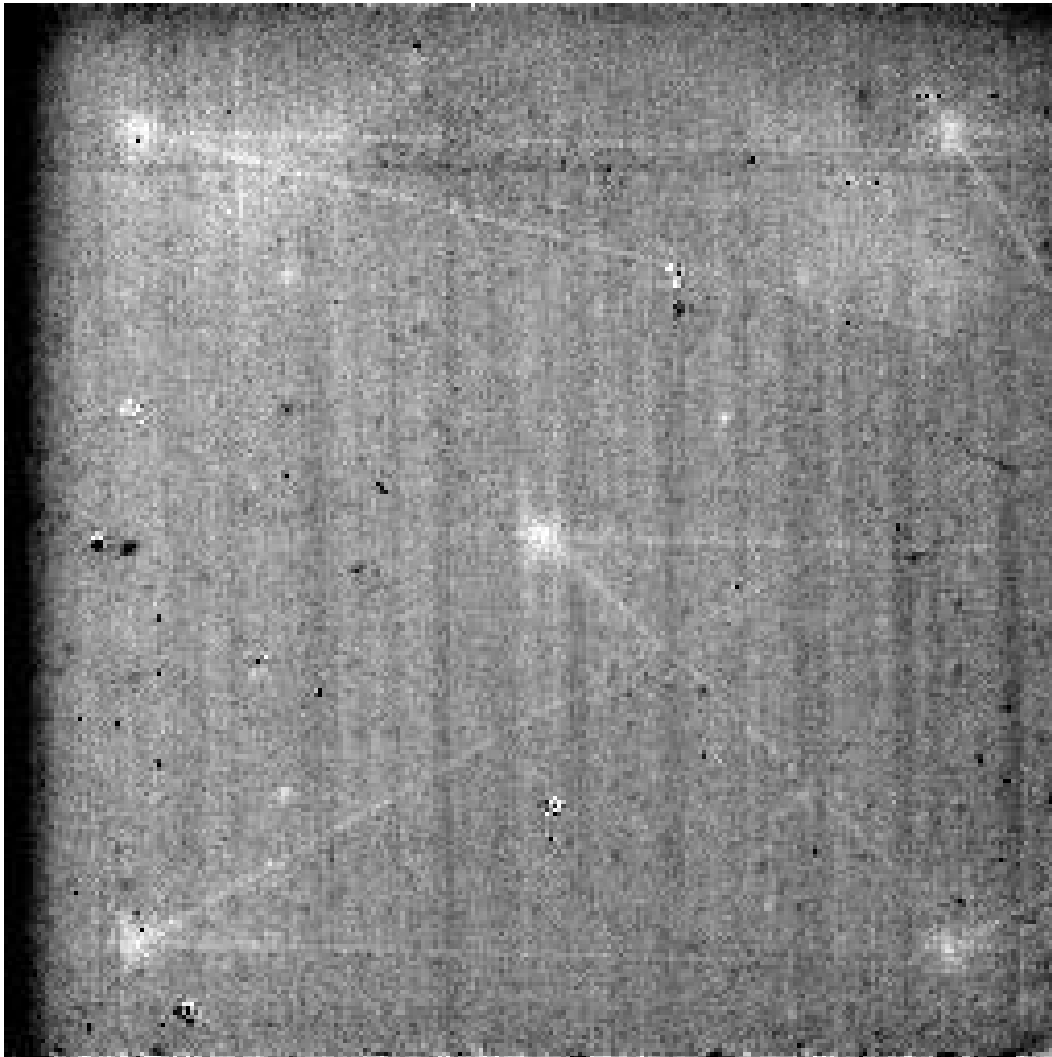
Figure 4.8: Median of channel 1 images from a calibration observation performed *after* observing Polaris. The 5 bright spots are persistent images from staring at the star while observing, while the set of criss-crossing lines were generated by slews between the various pointings. These observations were taken from AORKEY=3835904, in program pid=19.

Finally, we show an example of persistent images. Note that not all cases will be this obvious. In Figure 4.8 we see not only residual images of the star Polaris, but also latents left by Polaris as the telescope moved between dither positions. Sufficient dithering and using the SSC-provided outlier rejection provides effective mitigation of most residual images.

4.4 Optical Artifacts

4.4.1 Stray Light from Array Covers

Stray or scattered light on the arrays can be produced by illuminating regions off the edges of the arrays. Stray light from outside the IRAC fields of view is scattered into the active region of the IRAC detectors in all four channels. The problem is significantly worse in channels 1 and 2 than in channels 3 and 4. Stray light has two implications for observers. First, patches of stray light

can show up as spurious sources in the images. Second, background light, when scattered into the arrays, is manifest as additions to the flatfields when they are derived from observations of the sky. The scattered light is an additive, not a multiplicative term, so this will result in incorrect photometry when the flat field is divided into the data unless the scattered light is removed from the flat. Stars which fall into those regions which scatter light into the detectors produce distinctive patterns of scattered light on the array. We have identified scattered light avoidance zones in each channel where observers should avoid placing bright stars if their observations are sensitive to scattered light.

Figure 4.9 shows the zones for channels 1 and 2 with Spot overlays. Zones 1A, 1B, 2A and 2B (which produce the strongest scattered light) typically scatter about 2% of the light from a star into a scattered light "splatter pattern" which has a peak of about 0.2% of the peak of the star. Figures 4.10 to 4.13 show examples of stray light in channels 1–4. More examples and details on the locations of the stray light zones are on the SSC web pages at http://ssc.spitzer.caltech.edu/irac/scatt.html.
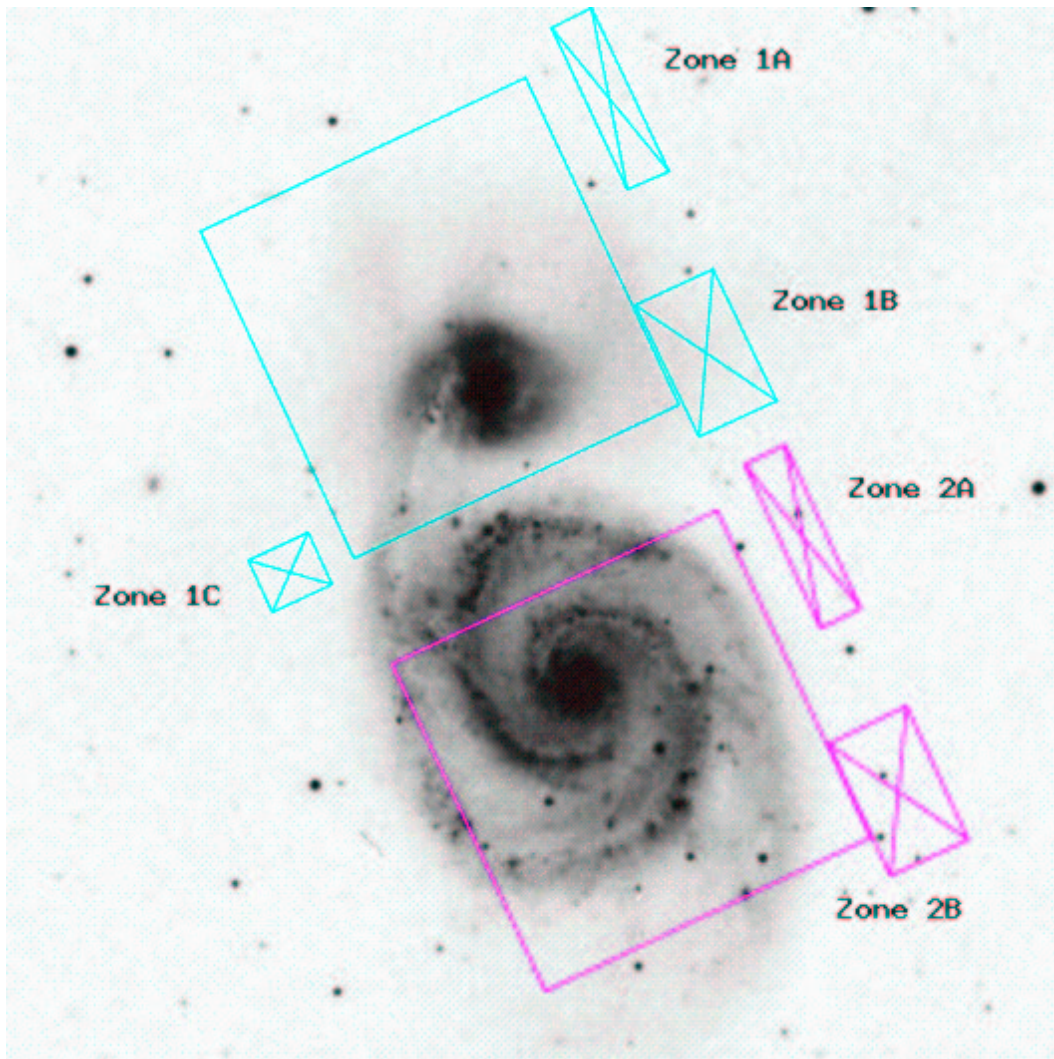

Figure 4.9: An image of the M51 system, showing an overlay of the IRAC fields of view, with the scattered light origin zones for channels 1 and 2 labeled.

Both point sources and the diffuse background generate stray or scattered light. Stray light due to the diffuse background is removed in the pipeline by assuming the source of illumination is

30

uniform and has a brightness equal to the COBE/DIRBE zodiacal light model. This assumption is not true at low Galactic latitudes or through interstellar clouds, but in the 3.6–8 μm wavelength range it is nearly correct. A scaled stray light template is subtracted from each image, in both the science and calibration pipelines. Before this correction was implemented, diffuse stray light from scattered zodiacal background contaminated the flats, which are derived from observations of high zodiacal background fields, and led to false photometric variations of 5%–10% in the portions of the array affected by stray light; this photometric error is now estimated to be less than 2%.

Example images of scattered light are shown here to alert you in case you see something similar in your IRAC images. The scattered light pattern from point sources is difficult to predict, and very difficult to model for removal. To the first order, you should not use data where scattered light from point sources is expected to cover or appears to cover your scientific target. Using sufficient redundancy (specifically, several well-spaced dither positions) in obtaining your data should still allow you to perform useful science with your images. A stray light masking IDL procedure is available from the SSC contributed software web page. This procedure incorporates our best understanding of the stray light producing regions. The procedure updates the corresponding imask or dmask for a BCD by determining whether a sufficiently bright star is in a stray light producing region. The 2MASS point source list, supplied as a post-BCD pipeline data product, is used to determine the bright star positions.

Figures 4.10 to 4.13 are 201 pixels (4.'1) square, and have been extracted from larger mosaics produced from the IRAC GTO shallow survey (PROGID 30). This survey covers 9 square degrees with three 30-second images at each position. Because the mosaics cover large areas, the star causing the scattered light appears in many of the images. All of the sample images have the same array orientation as the BCD images. The sample images are mosaics of a BCD that contains the stray light and the BCD that contains the star that produces the stray light. Figures 4.10 and 4.11 show scattered light in the two short-wavelength channels, from zones 1A, 1B, 2A, and 2B. Figures 4.12 and 4.13 show examples of the scattered light in channels 3 and 4. Because stars are much fainter in these channels, and the scattering geometry is much less favorable, these scattered light spots are much less obvious than in the short-wavelength channels. Dithering by more than a few pixels will take the bright star off the channel 3 and 4 "scattering strip," so the scattered light spots should be removed from mosaics made with adequately-dithered data.

Please note that Figures 4.10 to 4.13 were made with *no outlier rejection*. A dithered observation, combined with outlier rejection, will have much reduced stray light. Further, a diligent data analyst who recognizes and masks stray light in the individual BCDs will be able to eliminate stray light from well-planned mosaics. Observations made with little or no redundancy, or with dithers on scales smaller than the size of the stray light patches, will contain stray light and should be used with caution.
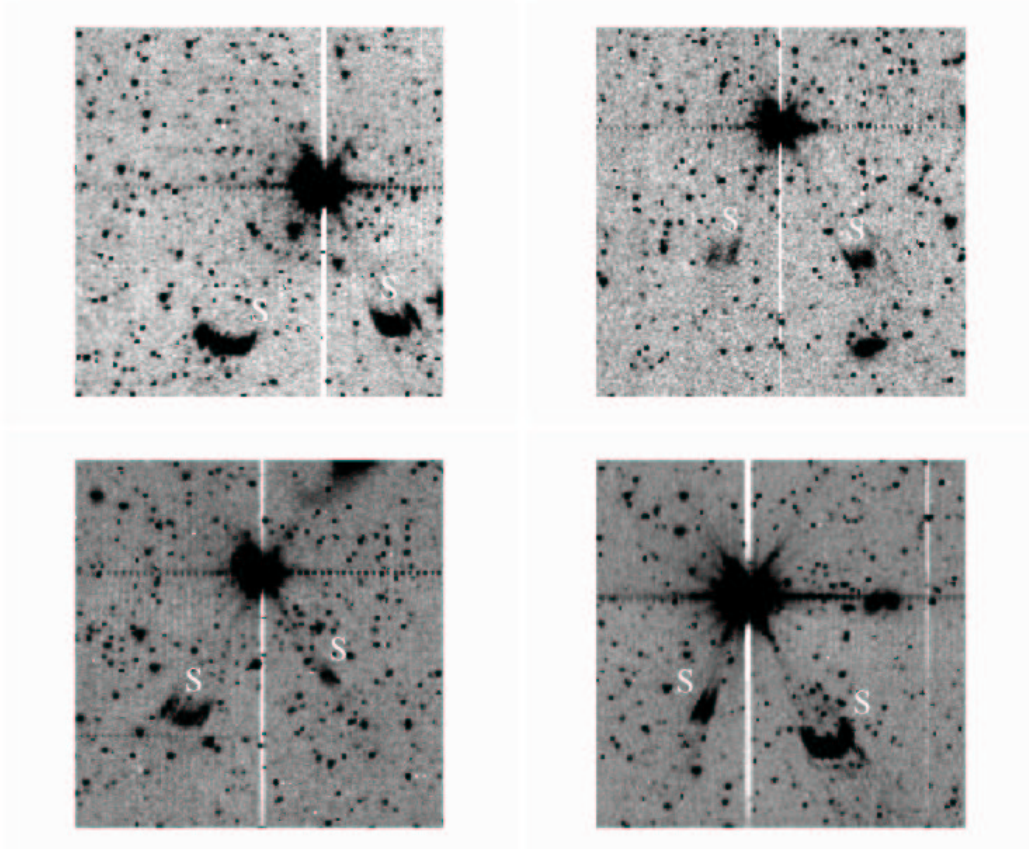
Figure 4.10: Channel 1 image showing scattered light on both sides of a bright star. The scattered light patches are marked with white "S" letters. The images were taken from program PID 30 data.
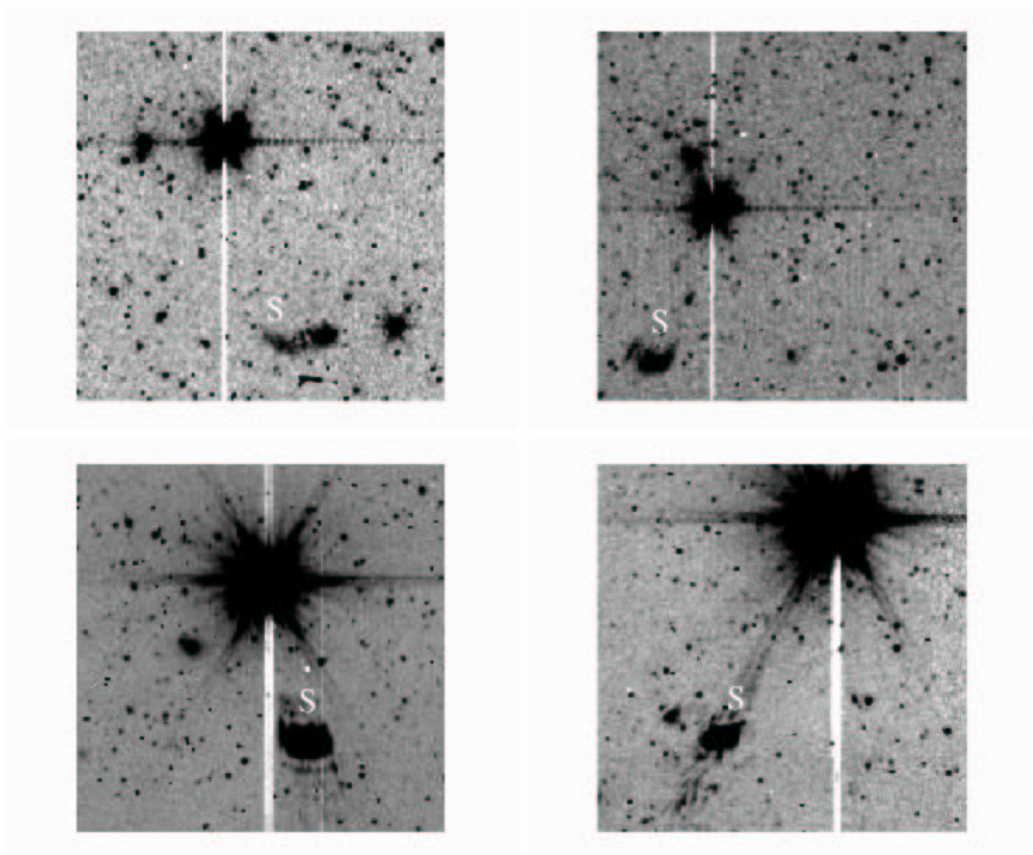
Figure 4.11: Channel 2 image showing scattered light on one side of a bright star. The scattered light patches are marked with white "S" letters. The images were taken from program PID 30 data.
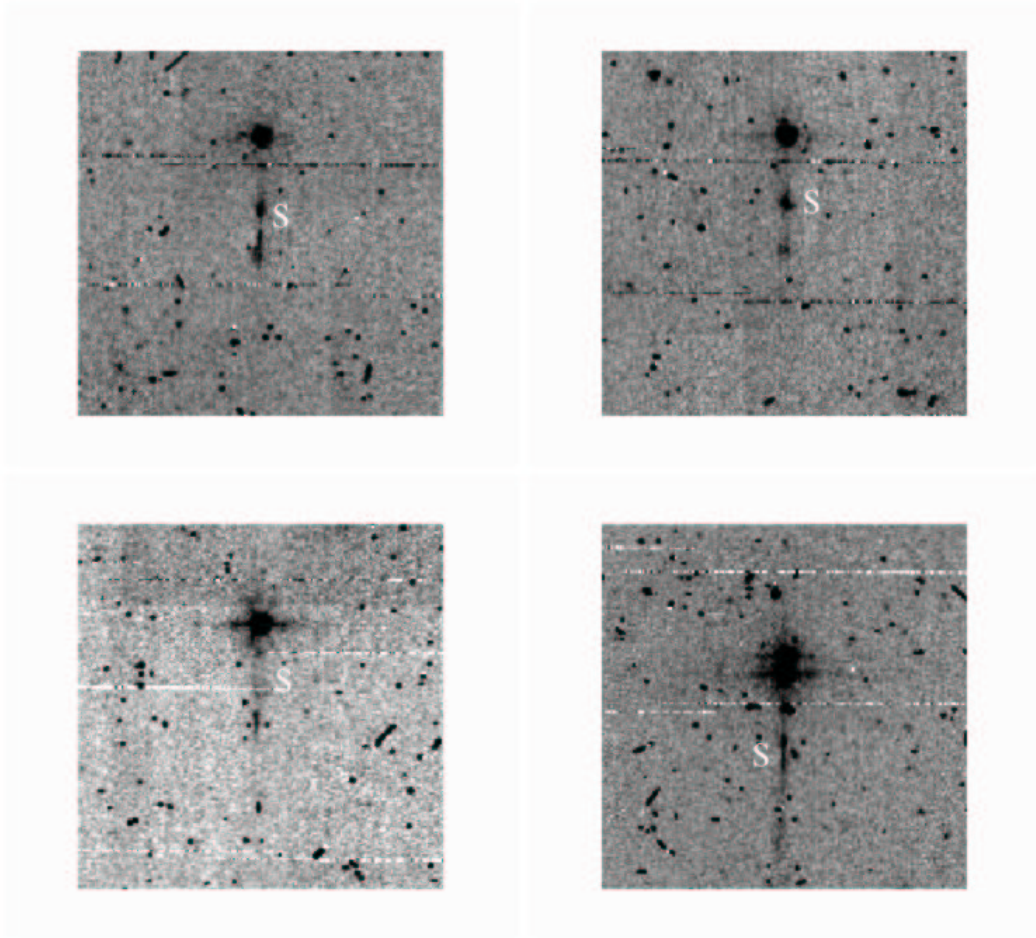
Figure 4.12: Channel 3 image showing scattered light from a scattering strip around the edge of the array where a bright star is located. The scattered light patches are marked with white "S" letters. The images were taken from program PID 30 data.
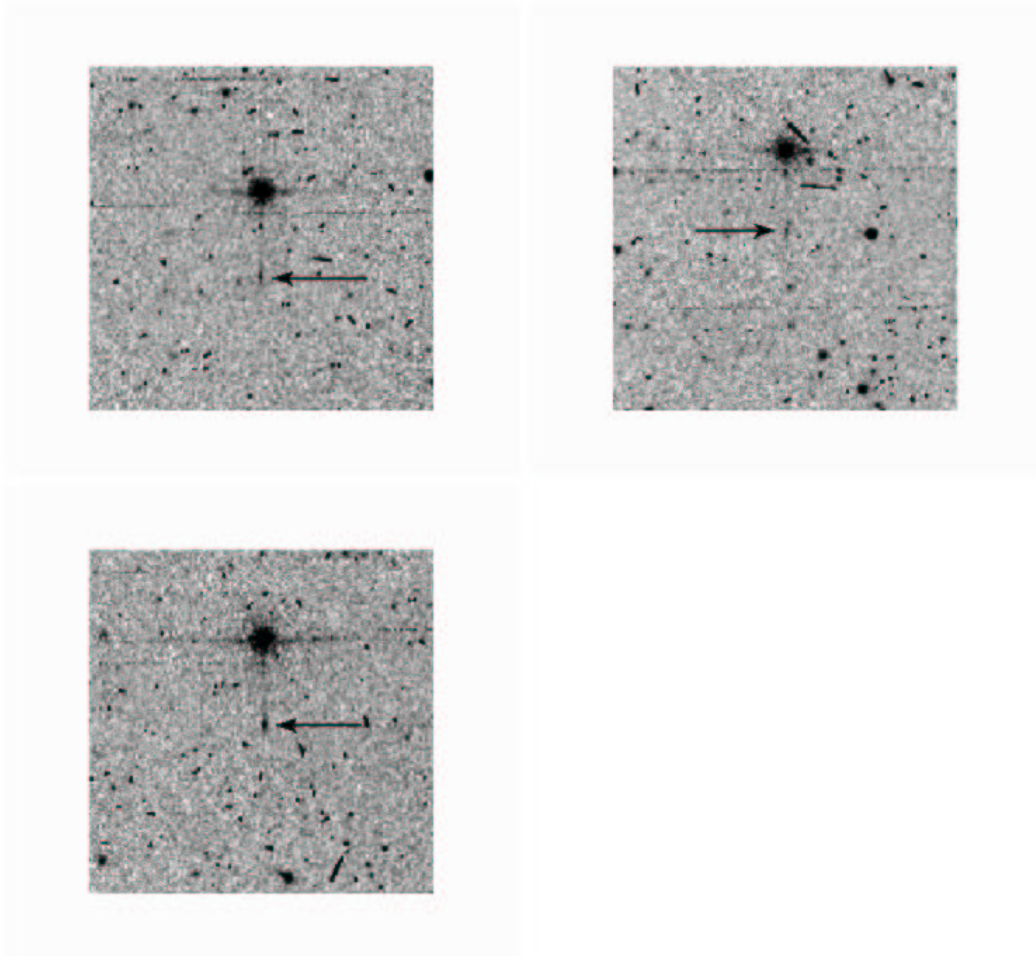
Figure 4.13: Channel 4 image showing scattered light from a scattering strip around the edge of the array where a bright star is located. The scattered light patches are pointed to by black arrows. The images were taken from program PID 30 data.

4.4.2   Optical Banding and Internal Scattering

The banding effect manifests itself as the rows and columns that contain a bright source having an enhanced level of brightness. This happens only in the Si:As arrays and has been shown to be due to internal optical scattering (inside the array). Both bright stellar sources and bright extended sources cause banding. It is clearly different from the optical diffraction patterns and the column pull-down effect. The SSC pipeline currently does nothing with regard to banding.
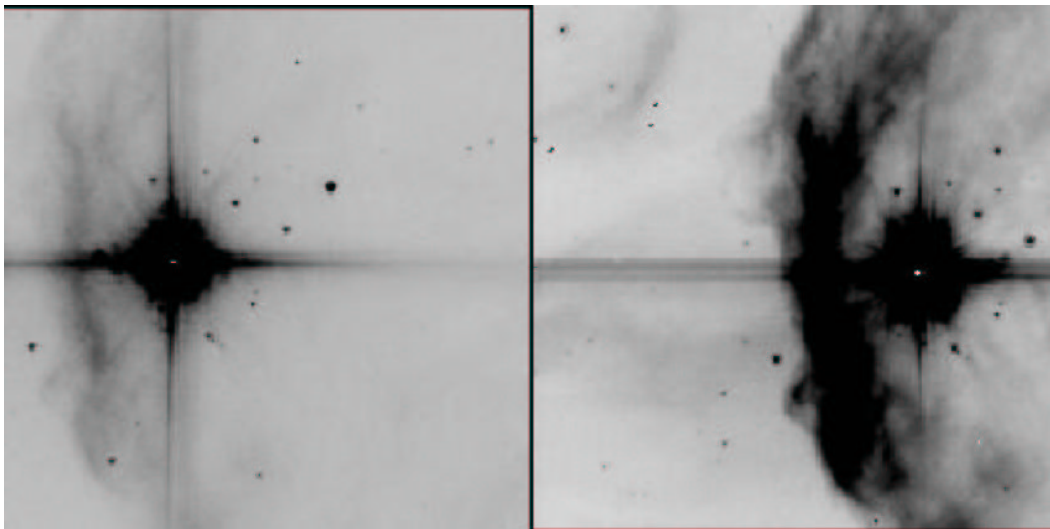
Figure 4.14: Typical image sections showing the banding effect. These are channel 3 (left) and channel 4 images of the same object (S140), adopted from a report by R. Gutermuth. These data were taken from program pid = 1046, AORKEY = 6624768.

Banding only appears in IRAC channels 3 and 4 (5.8 and 8 micron bands), and it is stronger in channel 3. Banding probably occurs at all intensity levels, but only appears obvious around bright sources that are at or near saturation levels. Banding is seen both in row and column directions, though their relative intensities are somewhat different. In addition, there is an electronic effect. Channel 4 has a strong row pull-up, and channel 3 has a weak column pull-up. The column pull-up is uniform across the row where the source is bright. The optical banding intensity falls off with distance from the bright spot. Cosmic ray hits cause electronic banding, but not optical banding.

The optical banding is only an enhancement of the optical scattering in channels 3 and 4 near the row and column where the source is. Approximately 25% of the light incident from a point source is scattered throughout the channel 3 array. The detected scattered light falls with distance from the source. Channel 4 has the same problem to a smaller degree. Recent laboratory tests have confirmed the large amount of optical scattering within the Si:As arrays. We are continuing to study this effect in the laboratory.

Users should be aware of the uncertainties resulting from banding, specifically when attempting measurements of faint sources near the affected rows or columns. For bright sources with significant banding, aperture photometry may not be successful, and it would be better to measure these sources using frames of shorter exposure times. Users are encouraged to experiment with image restoration techniques of their choice and discuss the relative merits of their approach with the IRAC team. Algorithms similar to the pulldown corrector may have some effectiveness in mitigating banding. The electronic artifact corrector on the SSC contributed software website will mitigate optical banding but does not model the flaring of banding towards the edges of the array.

Figure 4.15: Filter and beamsplitter ghosts.

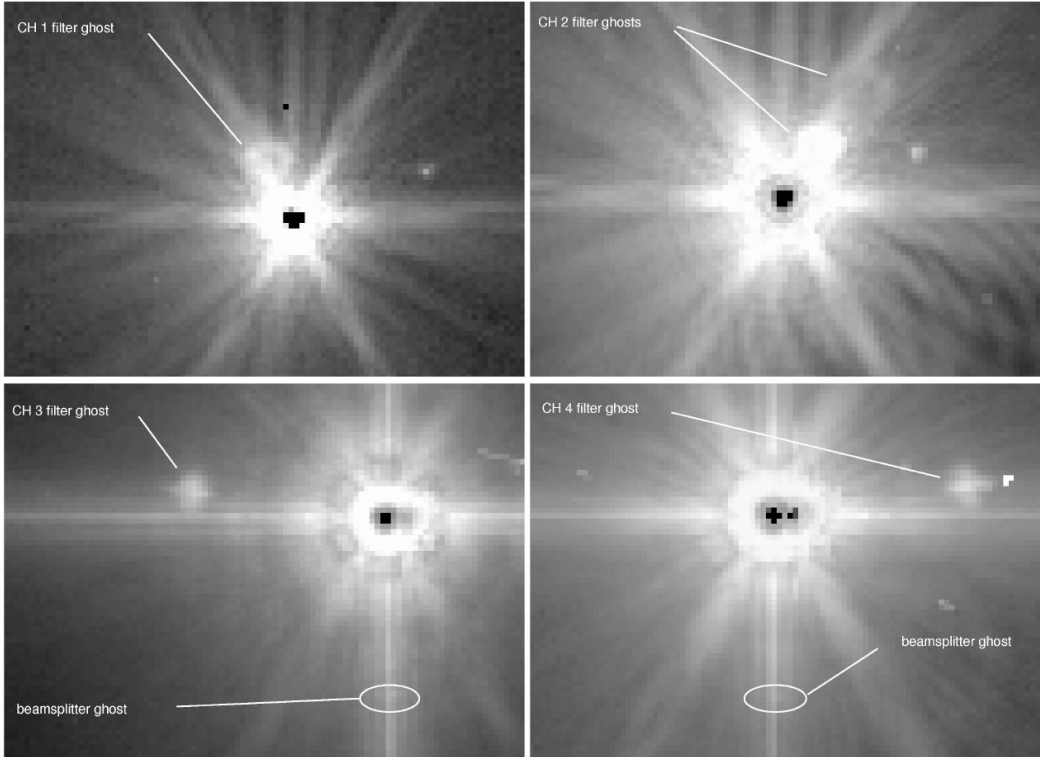### 4.4.3 Optical Ghosts

There are three types of known or potential optical ghosts visible in IRAC images. The brightest and most common ghosts are produced by internal reflections within the filters. The first-order filter ghosts (one pair of internal reflections) in channels 1 and 2 are triangular, and in the BCD images they appear above and/or to the left of the star in channel 1, and above and/or to the right of the star in channel 2. The channel 1 first order filter ghost contains $\sim 0.5\%$ of the flux of the main PSF in channel 1, and the channel 2 ghost $\sim 0.8\%$ of the flux of the channel 2 PSF. The second-order filter ghosts (two pairs of internal reflections) are much fainter ($\approx 25\%$ of the flux and $\approx 6\%$ of the surface brightness of the first order ghosts), rounder, larger, and about twice as far away from the star. The separation between the star and its ghosts increases with distance from the optical axis of the telescope. The channel 3 and 4 filter ghosts appear as small crosses at a larger distance, mostly to the left or right of the star, respectively, and contain $\lesssim 0.2\%$ of the flux of the main PSF in these channels. The separation and orientation are different from channels 1 and 2 because of the different orientations of the filters. The locations of the first-order filter ghosts can be estimated using the formulas given in Section 6.1.2.2.3 of the Spitzer Observer's Manual. Examples of filter ghosts are shown in Figure 4.15.

Similar ghosts are created by internal reflections within the beamsplitters. These only affect channels 3 and 4 which are transmitted through the beamsplitters. They appear as a very faint, short, horizontal bar at $\Delta y \approx -36$ pixels relative to a bright star (Fig. 4.15), but are often obscured by brighter "banding" artifacts. They are slightly fainter than the filter ghosts.

The faintest identified ghosts appear as images of the Spitzer entrance pupil, i.e., the primary mirror shadowed by the secondary and supports. These pupil ghosts are only found in channels 2 and 4, and require an extremely bright source (e.g., a first magnitude star in a 12-second frame) to be seen due to their low surface brightnesses. The pupil image is at a fixed location on the array (but in different locations in different channels). However, the pupil image is only partially
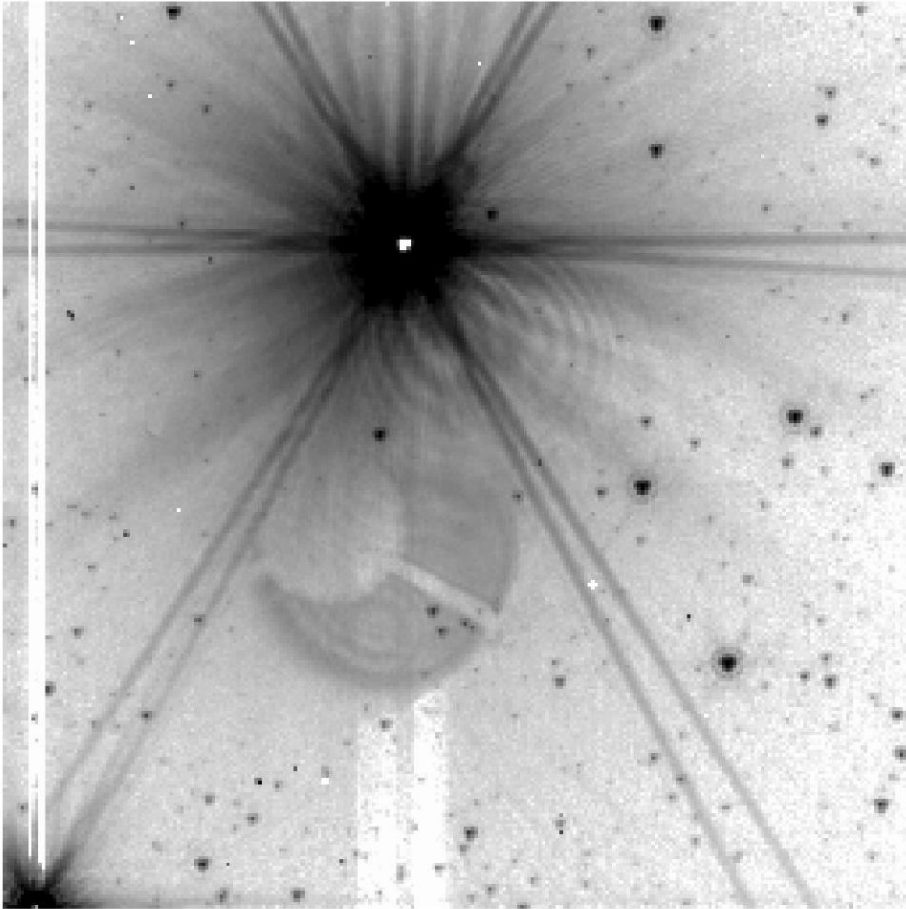
Figure 4.16: Pupil ghost in channel 2 from V416 Lac.

illuminated by a single source, and the portion of the image that is illuminated depends on the source position on the array. An example of a pupil ghost is shown in Figure 4.16. This figure also shows some fringes in between the pupil ghost and the star. At this time it is not clear if the fringes are directly related to the ghost. The total flux in these ghosts is $\sim 0.05\% - 0.5\%$ of the total flux in the PSF.

Currently we have no model for the exact shape and brightnesses of the ghosts, but we expect to develop models in the future. However, because the relative locations of the ghosts do vary with position on the array, sufficiently large dithering can help reduce or eliminate their effects. The stray light masking software also will flag the filter ghosts. The PRFs currently provided by the SSC include all ghosts, and the apertures used in calibrating channels 1 and 2 include the filter ghosts. In performing photometry for channels 1 and 2, the filter ghosts should be included.

4.4.4 Large Stray Light Ring and Splotches

There is a faint ring of scattered light with a mean radius of about 23.7 arcminutes in channel 1 that is visible around bright objects. There is also a slightly larger, fainter ring in channel 2. They were first noticed in mosaics of a SWIRE field that was adjacent to the bright star Mira in PID 181 (see Fig. 4.17). The ring was visible in channel 1 and 2 mosaics. We verified that the ring was an artifact by observing the field near Beta Gru where we observed pieces of the ring in the same places relative to the star. Starlight that is specularly reflected off the telescope mirrors cannot enter the MIC directly when the star is more than 16 arcminutes off the telescope boresight. IRAC's pupil stop is a little oversized, so the ring is probably light that is once or twice diffusely scattered at areas outside the secondary and/or near the top of the primary conical baffle. The mean surface brightness of the ring in channel 1 is 4.5×10^{-10} ($\pm$ 30%) times the mean surface brightness in the center pixel of a pixel-centered point source. Presumably there are stray light rings in channels 3 and 4, but they are too faint to see.

The "splotches" (see Fig. 4.18) are areas of more concentrated stray light that appear when a bright source is about 20–32 arcminutes off the center of the FOV, to the left or right in array coordinates. The splotches were seen in the SWIRE field near Mira and in the Beta Gru tests in channels 1, 2 and 4. The presence or absence of a splotch is very sensitive to the position relative to the bright object – a bright splotch can be present in one image and absent in an image with the telescope pointed a few pixels away. Even fainter splotches appeared in channel 2 about 1 degree away from Beta Gru, along the same directions.

Details of the large-angle stray light phenomena will be published in a paper by M. Lacy, R. Arendt, W. Glaccum, and J. Hora that is in preparation. We are investigating to find out where the scattering occurs that gives rise to the ring and the splotches. We thank R. Arendt for providing us with most of the information that was presented in this section.

4.5 Cosmic Rays and Solar Protons

The SSC mosaicker, `mopex`, identifies energetic particle hits as follows. All pixels in BCDs that contribute to a given pixel in the final mosaic are identified, and significant outliers (a user-specified number of sigmas above or below the filtered mean of all overlapping pixels of overlapping BCDs) are rejected. This method is very similar to the outlier rejection performed by shifting and adding ground based images. The rejected pixels can be inspected in the "Rmask" output files (one per input image). Outlier rejection in `mopex` can be adjusted. The parameters used in the online pipeline-generated mosaics rely on multiple (≥ 3) sightings of each sky pixel. In general, a coverage of at least five is necessary to produce optimal results with the multi-frame (standard) outlier rejection.
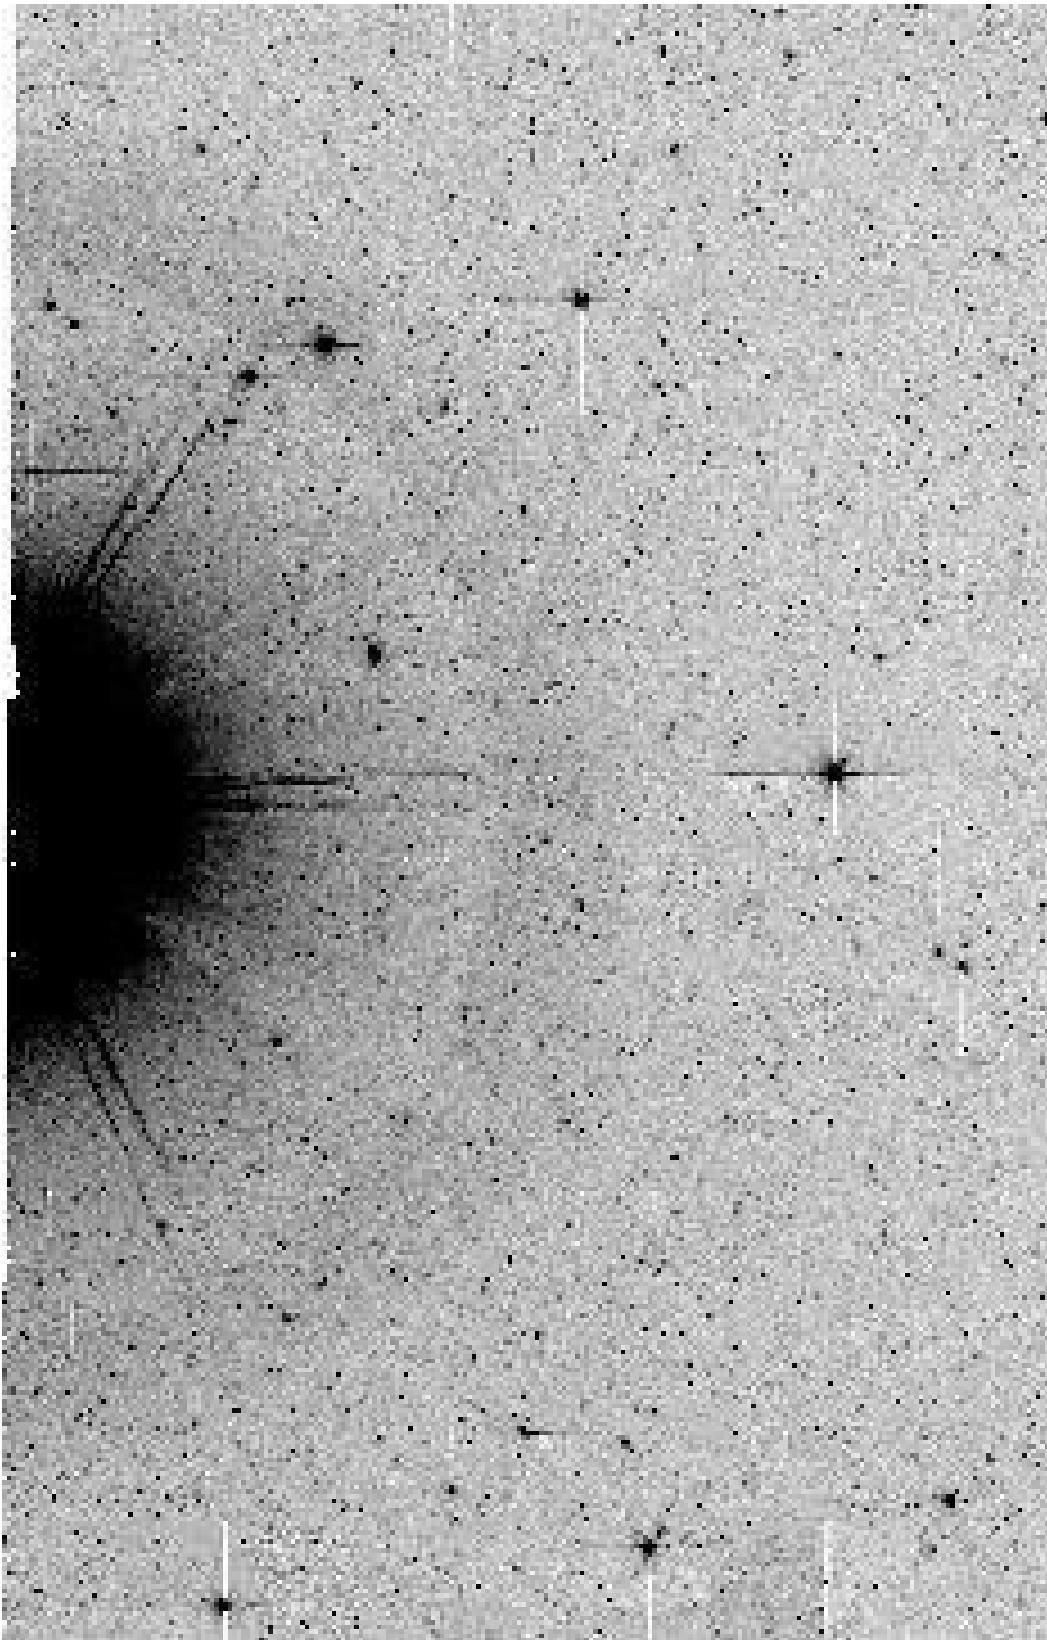
Figure 4.17: Part of the channel 1 mosaic (from observations in PID 181; AORKEYs 5838336, 5838592, 5839872 and 5840128) of the SWIRE field near Mira showing the 24 arcminute radius ring of stray light from the telescope.
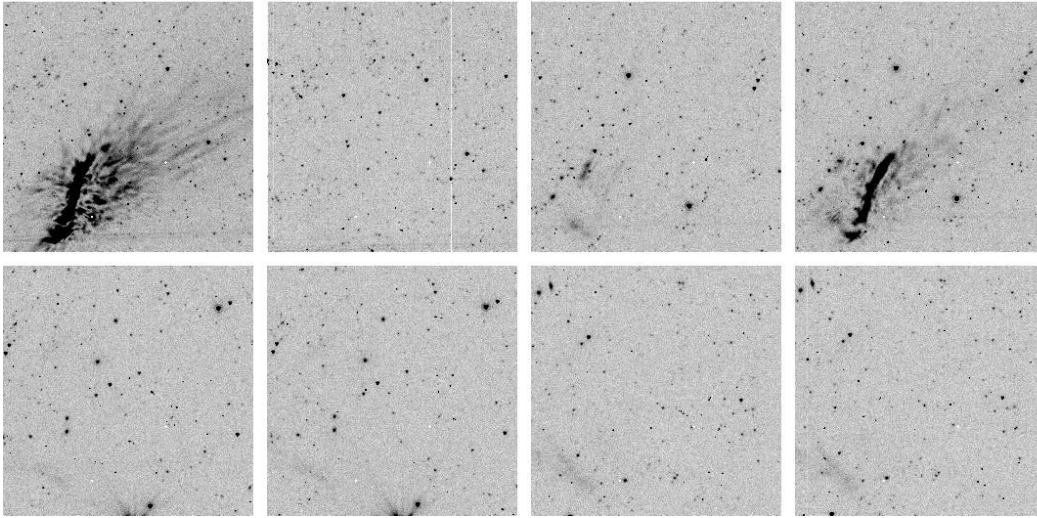
Figure 4.18: Channel 2 images from the SWIRE map showing stray light splotches from Mira, which was about 30 arcminutes away. Successive pairs of images were slightly dithered. The last pair is about 5 arcminutes from the first pair, but has a similar splotch. Note the absence of any stray light in the second image, though it was centered only a few pixels away from the first image. The images are from PID 181, AORKEY 5838336; EXPID 187-192, 199, and 200.

A special outlier rejection scheme can be used for sparse $(2-4\times)$ coverage; this "dual outlier" mode can be turned on using the namelist parameter file. Dual outlier rejection identifies pixels greater than a specified threshold above the background, groups these pixels and adjacent pixels above a threshold into objects, and compares the object to objects in overlapping frames. If the object overlaps with objects in other frames (in celestial coordinates), then it is not a cosmic ray. If the object is not detected in a user-specified fraction of overlapping images, it is flagged as a cosmic ray. This information is written into the Rmask files used for mosaicing and source extraction. The dual outlier method should also be used in conjunction with the multi-frame outlier rejection method. Multi-frame rejection may throw out data around bright sources depending on the thresholding, due to pixel phase effects between BCDs. Using the dual outlier rejection and the REFINE_OUTLIER=1 option will prevent this.

Additionally, a single-frame radiation hit detector is run and produces bit 9 in the Dmask, but this bit is not used by the SSC post-BCD software and is not recommended because radiation hits cannot be uniquely separated from real sources in single images.

Cosmic rays for channels 3 and 4 are larger and affect more pixels than the channel 1 and 2 cosmic rays due to the larger width of the active layer of the Si:As detectors. Some tuning of cosmic ray detection parameters may be necessary when working with deep integrations, especially for channels 3 and 4.

Each IRAC array receives approximately 1.5 cosmic ray hits per second, with ~ 2 pixels affected in channels 1 and 2, and ~ 6 pixels per hit affected in channels 3 and 4. The cosmic ray flux varies randomly by up to a factor of a few over time scales of minutes but does not undergo increases larger than that. Also, the cosmic ray flux is normally about a factor of two **higher** on average around solar **minimum** compared with solar maximum. The next solar minimum is expected to occur around the end of 2006.

Radiation hits do increase suddenly and dramatically during some major solar proton events. Historically, several such events have occurred over the course of the active part of a solar cycle. However, Spitzer might possibly have seen its last major proton event because we are approaching the solar cycle minimum when such events are normally very unlikely.
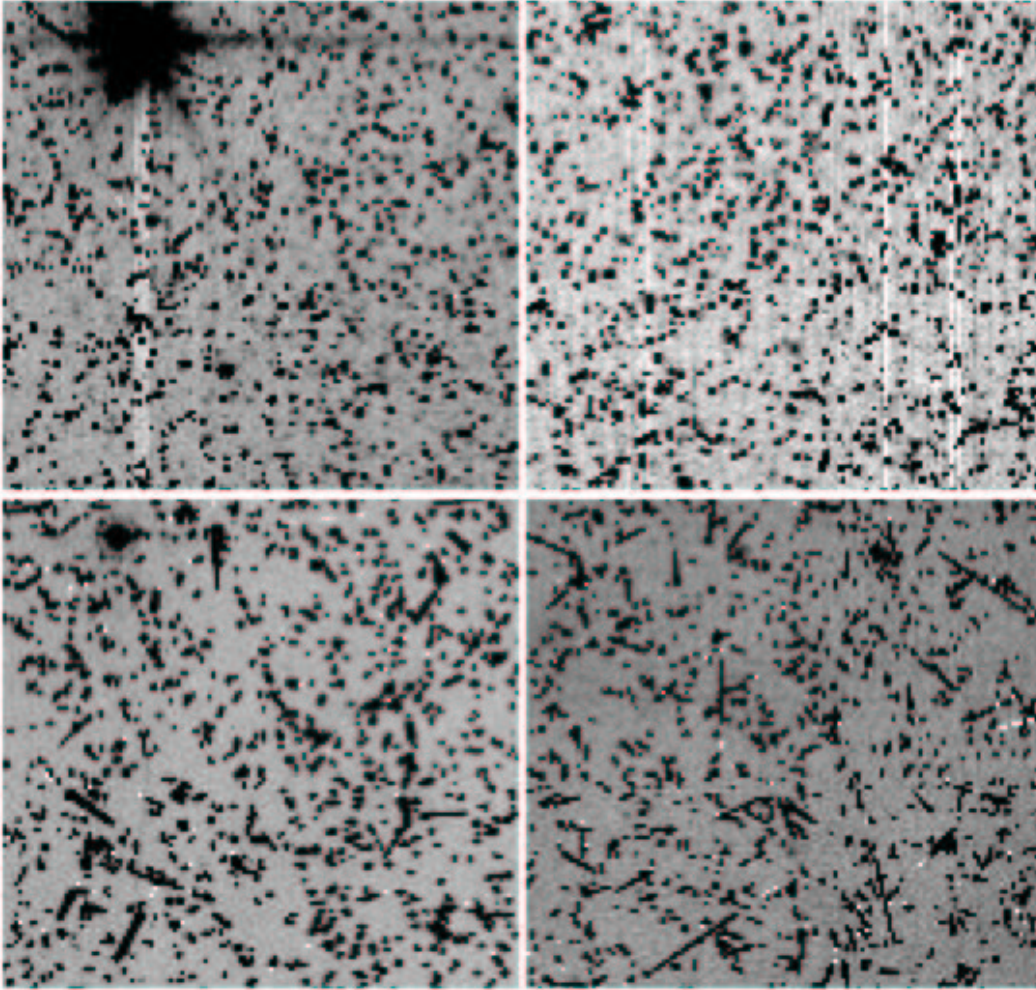
Figure 4.19: The central 128×128 pixels of IRAC 12-second images taken on January 20, 2005 during a major solar proton event. Channels 1 and 2 are top left and top right; channels 3 and 4 are bottom left and bottom right. Except for the bright star in channels 1 and 3, almost every other source in these images is a cosmic ray. These data are from observations in pid 3126.

Two major solar proton events occurred during IOC, so the SSC has experience in identifying them and their effects. Because of shielding around the instruments, only extremely energetic protons (> 100 MeV) of any origin appear as cosmic ray hits in the data. Thus, many solar weather phenomena ("storms," etc.) which do occasionally affect other spacecraft, or ground systems, are not of concern to *Spitzer*.

The Space Environment Center maintains a web site which reports near real time space weather, using data from GOES spacecraft in Earth orbit. However, because *Spitzer* is receding from the Earth at the rate of about 0.1 AU per year, or approximately 15 degrees annually along the Earth's orbital path, as the mission elapses GOES data become increasingly uncertain measures of the particle environment at *Spitzer*.

Observers with severely contaminated data (preventing them from reaching their proposed scientific goals) can request (via the Helpdesk) that their observation be remade (submit request within two months of receiving access to your proprietary data).

Radiation has very little effect on the IRAC arrays beyond elevating the counts in a given pixel. Repeated bombardment will eventually result in elevated dark current levels; however, the only effect to date is a slight increase in moderately hot pixels in channel 4 and no significant degradation of the arrays is expected throughout the observatory lifetime. Anneal procedures (heating and cooling of the arrays) are carried out at the end of every Period of Autonomous Operations (PAO) to restore the performance of the arrays. After anneals, all calibrations are repeated.

Some high energy cosmic rays cause persistent images, column pull-down, and muxbleed effects.

Chapter 5

Photometry and Calibration

5.1 Absolute Calibration

5.1.1 Current IRAC Absolute Calibration

IRAC is calibrated using both so-called primary and secondary calibrator stars. The primary stars are used to monitor long-term variations in the absolute calibration. They currently number 11 stars, are located in the continuous viewing zone (CVZ), and are thus observable year-round. They are observed once at the beginning, and once at the end of each campaign, i.e., about every 10 days whenever the instrument is switched on. The primary calibrators (increasing magnitude) are (J2000):

```
NPM1+67.0536 = SAO 17718 = 2MASS J17585466+6747368 (K2III, Ks=6.4)
HD 165459 = 2MASS J18023073+5837381 (A1V, Ks=6.6),
NPM1+68.0422 = BD+68 1022 = 2MASS J18471980+6816448 (K2III, Ks=6.8)
KF09T1 = GSC 04212-01074 = 2MASS J17592304+6602561 (K0III, Ks=8.1),
NPM1+66.0578 = GSC 04229-01455 = 2MASS J19253220+6647381 (K1III, Ks=8.3),
NPM1+64.0581  = HD 180609 = 2MASS J19124720+6410373 (A0V, Ks=9.1),
NPM1+60.0581 = BD+60 1753 = 2MASS J17245227+6025508 (A1V, Ks=9.6),
KF06T1 = 2MASS J17575849+6652293 (K1.5III, Ks=11.0),
KF08T3 = 2MASS J17551622+6610116 (K0.5III, Ks=11.1),
KF06T2 = 2MASS J17583798+6646522 (K1.5III, Ks=11.3),
2MASS J18120956+6329423 (A3V, Ks=11.6)
```

The secondary calibrator stars are used to monitor short-term variations in the absolute calibration. To avoid slew overheads, they are observed close to downlinks and must therefore be located near the ecliptic plane, in a tightly constrained window of about 20 degrees. Because of the motion of the Earth about the Sun this window constantly moves and so any one secondary calibrator is visible for only a campaign or two per year. In practice, the calibration values for IRAC appear to be quite temporally stable.

The data are calibrated by means of aperture photometry, using a 10 native pixel radius (12 arcsec) aperture. Unfortunately, ground-based infrared calibrators are too bright to use as calibrators for IRAC. Therefore, one must use models to predict the actual flux for each channel as a function of spectral type (Cohen et al. 2003). Table 5.1 lists the calibration factors that are used in the S13 pipeline. The absolute calibration is described in detail in Reach et al. (2005), with further refinements at the 1–3% level.

The absolute gain calibration is accurate to better than 3%. The stellar photometry is repeatable at the $< 1\%$ level. The absolute fluxes of the calibration stars are known to 2%–3% (Cohen et al. 2003). To obtain photometry at this accuracy, photometric corrections for the

Table 5.1: The photometric calibration and zero magnitude flux for IRAC

λ	Calfac	FLUXCONV	F_{ν_0}
(μm)	(DN/s)/μJy	(MJy/sr)/(DN/s)	(Jy)
3.6	0.255	0.1088	280.9$\pm$4.1
4.5	0.207	0.1388	179.7$\pm$2.6
5.8	0.046	0.5952	115.0$\pm$1.7
8.0	0.133	0.2021	64.1$\pm$0.9

location of the source within its peak pixel, and the location of the source within the array[1], must be made.

Note that IRAC is not an absolute background photometer, so the total brightness in IRAC images should be used with great caution. Absolute background photometry is still under active investigation. There is a cold shutter in the calibration assembly, but it is not being operated in flight, in order to minimize mission risk. Therefore, the offset level in IRAC images is referred to laboratory measurements before launch, where the offset level was observed to change very significantly from one laboratory experiment to another.

In laboratory tests, the absolute offset of IRAC images was found to vary at levels that are comparable to the minimum celestial background in channels 1 and 2. Furthermore, the offset level changes depending on whether the detector has been recently annealed. Thus, for diffuse surface brightness measurements, we recommend making *differential* measurements among at least two sky positions, preferably from the same campaign.

5.1.2 Absolute Calibration History

At the time of writing of the current version of the Data Handbook, all IRAC data are being reprocessed using the S13 pipeline. Note that we generally tie the absolute calibration factors to pipeline versions, because they must be derived self-consistently between the calibration pipelines (such as those that generate the campaign darks and flats) and science pipeline (which applies darks, flats, and linearity). The absolute calibration factors (FLUXCONV) used for the S13 reprocessing are documented in a refereed journal article (Reach et al. 2005). We validated these calibration factors by reprocessing a campaign's worth of primary and secondary calibration stars using S13 and the updated calibration files (improved linearity curve for channel 3, superflat from two years worth of flat fields). Results indicated that the stellar fluxes derived from the data products in electrons/sec, which are independent of FLUXCONV, are identical between S11, S12 and S13. We therefore updated the absolute calibration factors to match the PASP paper.

The previous absolute calibration factors, used for pipeline versions S11 and S12, were derived using the same principles but a preliminary version of the analysis. The only notable change between the S11/S12 and S13 calibrations is a 3% change in IRAC channel 4. This is well within the 5% absolute calibration uncertainty that we recommend to be cited in papers. Careful analysis performed according to the prescription in the PASP paper can achieve higher relative and absolute accuracy, but we recommend citing 5% uncertainty for general observers, assuming that all the proper corrections (see the following sections) have been made.

---

[1]http://ssc.spitzer.caltech.edu/irac/locationcolor/

5.2 Color Corrections

IRAC is a broad-band photometer. We describe here the method used for calibrating our data in specific surface brightness (MJy/sr) or flux density (Jy) units, and we provide the prescription for how to interpret the data for sources with spectral shapes other than the nominal one assumed in the calibration process.

The conventions used by IRAC are the same as those used by *IRAS* (Beichman 1988, IRAS Explanatory Supplement, §VI.C.3), *COBE*/DIRBE (Hauser 1998), and *ISO* (Blommaert et al. 1999). The basic idea is to quote a flux density $F_{\nu_0}^{\rm quot}$ at a nominal wavelength $\lambda_0 = c/\nu_0$ that would be accurate for a source with a nominal spectrum, $\nu F_{\nu}^{nom} =$constant. Using this F_{ν}^{nom} is merely a matter of convention; in fact a wide range of spectra are expected, both redder (e.g., interstellar medium, asteroids) and bluer (e.g., stars) than nominal. The color correction tables given below allow observers to convert the nominally-calibrated data, $F_{\nu_0}^{\rm quot}$, into more accurate estimates of the flux density at the nominal wavelength.

The number of electrons collected from the nominal source in a straight integration of duration t using a telescope with area A is

$$N_e^{nom} = tA \int \frac{F_{\nu}^{nom}}{h\nu} R d\nu,$$ (5.1)

where R is the system spectral response. The convention for R is that it is proportional to the number of electrons produced by a single photon with energy $h\nu$. If we define for convenience

$$\Delta\nu = \int (\nu/\nu_0)^{-2} R d\nu,$$ (5.2)

then the number of electrons collected from a source with the nominal spectrum is

$$N_e^{nom} = tA \frac{F_{\nu_0}^{nom}}{h\nu_0} \Delta\nu.$$ (5.3)

The calibration factor, by which the number of electrons, N_e from an arbitrary source must be divided in order to give the quoted flux density at the nominal wavelength, is

$$C = \frac{N_e^{nom}}{F_{\nu_0}^{nom}} = \frac{At\Delta\nu}{h\nu_0}.$$ (5.4)

The calibration factor is measured using observations of a celestial calibrator source of known spectrum, F_{ν}^*. The number of electrons collected from the star is

$$N_e^* = tA \int \frac{F_{\nu}^*}{h\nu} R d\nu.$$ (5.5)

Combining with equation 5.4, we can express the calibration factor in terms of the observed number of electrons from the calibration source:

$$C = \frac{N_e^* \Delta\nu}{h\nu_0 \int \frac{F_{\nu}^*}{h\nu} R d\nu}.$$ (5.6)

We can now cast this in convenient terms as follows:

$$C = \frac{N_e^*}{F_{\nu_0}^* K^*},$$ (5.7)

where $F_{\nu_0}^*$ is the flux density of the calibration source at the nominal wavelength, and K^* is the color correction factor for the calibration source spectrum.

The color correction factor for a source with spectrum F_ν is defined as

$$K = \frac{\int (F_\nu/F_{\nu_0}) (\nu/\nu_0)^{-1} Rd\nu}{\int (\nu/\nu_0)^{-2} Rd\nu}. \tag{5.8}$$

In this convention, the overall normalization of R is unimportant. Observers can correct the photometry to the spectrum of their source by either performing the integral in this equation or looking up the color corrections for sources with similar spectra. (Note that our definition of the color correction looks slightly different from that in the *IRAS Explanatory Supplement*, because we used the system spectral response R in electrons/photon, instead of erg/photon.)

We selected nominal wavelengths that minimize the need for color corrections, such that the quoted flux densities in SSC data products are minimally sensitive to the true shape of the source spectrum. (This paragraph can be skipped by most readers; the table is given below.) First, let us expand the source spectrum in a Taylor series about the nominal wavelength:

$$F_\nu = F_{\nu_0} \left[1 + \beta \left(\frac{\lambda - \lambda_0}{\lambda_0} \right) + ... \right]. \tag{5.9}$$

Using equation 5.8, the color correction for a source with spectrum F_ν is

$$K = \frac{1}{\Delta\nu} \int \left[1 + \beta \left(\frac{\lambda - \lambda_0}{\lambda_0} \right) \right] (\nu/\nu_0)^{-1} Rd\nu. \tag{5.10}$$

The choice of λ_0 that makes K minimally sensitive to β is the one for which

$$\frac{dK}{d\beta} = 0,$$

which is

$$\lambda_0 = \langle \lambda \rangle = \frac{\int \lambda (\nu/\nu_0)^{-1} Rd\nu}{\int (\nu/\nu_0)^{-1} Rd\nu} = c\frac{\int \nu^{-2} Rd\nu}{\int \nu^{-1} Rd\nu}. \tag{5.11}$$

So the optimum choice of λ_0 for insensitivity to spectral slope is the weighted average wavelength.

Table 5.2: IRAC nominal wavelengths and bandwidths.

Channel	λ_0 (μm)	R_{max}	Eff. Width $\Delta\nu$ (Eq. 5.2; 10^{12} Hz)	Width $\Delta\nu/R_{max}$ (10^{12} Hz)	1/2-power wavelengths blue	red (μm)
1	3.550	0.651	7.57	16.23	3.18	3.92
2	4.493	0.736	6.93	12.95	4.00	5.02
3	5.731	0.285	1.93	11.70	5.02	6.43
4	7.872	0.428	3.94	12.23	6.45	9.33

Using the nominal wavelengths from Table 5.2, the color corrections for a wide range of spectral shapes are less than 3%. Thus, when comparing IRAC fluxes to a theoretical model, placing the data points on the wavelength axis at λ_0 takes care of most of the potential color-dependence. To place the data points more accurately on the flux density axis, take the quoted flux densities derived from the images, and divide by the appropriate color correction factor in the tables below:

$$F_{\nu_0} = \frac{F_{\nu_0}^{quot}}{K} \tag{5.12}$$

Or, calculate the color correction using equation 5.8 together with the spectral response tables, which are available in the IRAC section of the SSC website.

Table 5.3 shows the color corrections for sources with power-law spectra, $F_\nu \propto \nu^\alpha$, and Table 5.4 shows the color corrections for blackbody spectra with a range of temperatures. The nominal spectrum has $\alpha = -1$, so the color corrections are unity by definition in that column. These calculations are accurate to $\sim 1\%$. Note that the color corrections for a ν^{-1} and a ν^0 spectrum are always unity. This is in fact a theorem that is easily proven using equations 5.8 and 5.11.

Table 5.3: Color corrections for power-law spectra, $F_\nu \propto \nu^\alpha$

	Color correction for $\alpha =$				
Band	-2	-1	0	1	2
1	1.0037	1	1	1.0037	1.0111
2	1.0040	1	1	1.0040	1.0121
3	1.0052	1	1	1.0052	1.0155
4	1.0111	1	1	1.0113	1.0337

Table 5.4: Color corrections for blackbody spectra

	Temperature (K)							
Channel	5000	2000	1500	1000	800	600	400	200
1	1.0063	0.9990	0.9959	0.9933	0.9953	1.0068	1.0614	1.5138
2	1.0080	1.0015	0.9983	0.9938	0.9927	0.9961	1.0240	1.2929
3	1.0114	1.0048	1.0012	0.9952	0.9921	0.9907	1.0042	1.1717
4	1.0269	1.0163	1.0112	1.0001	0.9928	0.9839	0.9818	1.1215

Table 5.5 gives the color corrections for the spectrum of the zodiacal light, which is the dominant diffuse background in the IRAC wavelength range. The first model is the COBE/DIRBE zodiacal light model as implemented in Spot. The zodiacal light is mostly due to thermal emission from grains at ~ 260 K over the IRAC wavelength range, except in channel 1 where scattering contributes $\sim 50\%$ of the brightness. The second zodiacal light spectrum in Table 5.5 is the ISOCAM CVF spectrum (5.6–15.9 μm; Reach et al. 2003) spliced with the COBE/DIRBE model at shorter wavelengths.

Table 5.5: Color corrections for zodiacal light spectrum

	COBE/DIRBE model			ISOCAM+COBE/DIRBE		
Band	I_{ν_0} (MJy/sr)	K	$I_{\nu_0}^{quot}$	I_{ν_0} (MJy/sr)	K	$I_{\nu_0}^{quot}$
1	0.067	1.0355	0.069	0.40	1.0355	0.42
2	0.24	1.0835	0.26	1.44	1.0835	1.56
3	1.11	1.0518	1.16	6.64	1.0588	7.00
4	5.05	1.0135	5.12	25.9	1.0931	28.4

For sources with complicated spectral shape the color corrections can be significantly different from unity. The corrections are infinite in the case of a spectrum dominated by narrow lines, because there may be no flux precisely at the nominal wavelength, which only demonstrates that such sources should be treated differently from continuum-dominated sources. We calculated one illustrative example which may prove useful. The ISO SWS spectrum of NGC 7023 is dominated by PAH emission bands and a faint continuum over the IRAC wavelength range. Table 5.6 shows the color corrections using equation 5.8 and the ISO spectrum. The large value in channel

1 is due to the presence of the 3.28 μm PAH band, which dominates the in-band flux relative to the weak continuum at the nominal wavelength of 3.535 μm. Channel 2 is mostly continuum. Then channel 3 is dominated by a PAH band at 6.2 μm. Channel 4 has significant PAH band emission throughout, with prominent peaks at 7.7 and 8.6 μm. The values in this table can be used for comparison to IRAC colors of other sources by anti-color-correction, which gives the predicted colors for NGC 7023 in the same units as the SSC calibrated data: $F_{\nu_0}^{quot} = F_{\nu_0} \times K$, which is shown in the last column of Table 5.6. Thus PAH-dominated sources are expected to have $F_{\nu_0}^{quot}(8\mu\text{m})/F_{\nu_0}^{quot}(5.8\mu\text{m}) = 599/237 = 2.5$.

Table 5.6: Color corrections for NGC 7023 (PAH-dominated) spectrum

Band	F_{ν_0} (Jy)	K	$F_{\nu_0}^{quot}$
1	17.3	2.21	38.3
2	30.3	1.21	36.6
3	169	1.40	237
4	1021	0.59	599

For observations of sources dominated by spectral lines, the quoted flux densities should be converted into fluxes using

$$F = \frac{F_{\nu_0}^{quot} \Delta\nu \lambda}{R_l \lambda_0} \tag{5.13}$$

where λ is the wavelength of the spectral feature and R_l is the spectral response at that wavelength. Both λ_0 and effective width $\Delta\nu$ are in Table 5.2. The formalism used for continuum sources is inappropriate for spectral-line sources because it is likely that $F_{\nu_0} = 0$ and $K = \infty$. It is important that the normalization of R used to determine $\Delta\nu$ and R_l is the same. In Table 5.2, the column $\Delta\nu$ (effective width) was calculated with the same normalization of response function as on the SSC website so it is the appropriate one to use. The maximum response, R_{max} is also given in that table, so the fluxes of lines in heart of the waveband can be estimated by simply multiplying the quoted flux densities by $\Delta\nu/R_{max}$, which is listed in the table in the column "Width."

5.3 Array-Location-Dependent Photometric Corrections for Compact Sources with Stellar Spectral Slopes

IRAC Basic Calibrated Data (BCD) are corrected for pixel-to-pixel gain variations, a process commonly known as "flat fielding". The IRAC flat field is derived by imaging the high surface-brightness zodiacal background. There are two relevant points about this approach. First, the zodiacal background is extended and essentially uniform on IRAC size-scales, and thus uniformly fills the field of view. Second, the zodical background peaks redward of the IRAC filters. The vast majority of objects seen by IRAC are not like this. Many are compact, being either stars or background galaxies. Many have spectral energy distributions in the IRAC filters more closely resembling stars. Stars (and many galaxies) have color temperatures that are fairly high, and peak blueward of the IRAC filters. Generally speaking, for these objects the IRAC filters are well on the Rayleigh-Jeans side of the blackbody spectrum.

There are several consequences of these two points. IRAC has significant scattering, as well as spatial distortion. As a result, the extended and point source effective gains are slightly different. In addition, there is a variation in the effective filter bandpass as a function of angle of incidence, which in turn depends on the exact position of an object on the array (Quijada et al.

2004). As a result of this, while the flat field perfectly corrects the extended zodical background (or any extended object) with a similar spectral slope, it is incorrect for many other objects.

This effect has been directly measured. Objects were sampled at many different locations on the array, and their flux measured from the BCD images. The systematic variations in their measured fluxes were used to derive the corrections. The amplitude of this effect is sizeable. It may reach 10% peak-to-peak, depending on the detector array. This is larger than any other source of uncertainty in the IRAC calibration. For well-dithered data, experiments have shown that this effect tends to average out so that the amplitude of the effect is very small (less than 1%). However, depending on the exact details of the mapping and dithering, it is not uncommon to have small areas of data where the mean correction approaches the full 10%.

Correction images may be downloaded from *http://ssc.spitzer.caltech.edu/irac/locationcolor/*.

Users should note the following:

- The correction images are oriented so that they apply multiplicatively to the BCD images. Among other things, the channel 1 and 2 arrays are flipped around their vertical axis during the reduction by the BCD pipeline, hence these images cannot be directly applied to the raw data.

- The correction images are for compact, or point-like sources.

- The correction images are for a Rayleigh-Jeans (stellar, Vega-like) spectrum. Spectral indices differing from this will have different corrections. Generally, most IRAC objects have spectral slopes that are bracketed by the two extremes of the red zodiacal spectrum and the blue stellar spectrum, so the corrections will lie between zero and that in the correction image.

- Note that the existing flat-field flattens the zodiacal background. After correction, although the point sources may be correct, the background will no longer be flat.

- For photometry on mosaics, observers will have to generate their own correction images.

5.4 Pixel-Phase-Dependent Photometric Correction for Point Sources

The flux density of a point source measured from IRAC images depends on the exact location where the peak of the Point Response Function (PRF) falls on a pixel. This effect is due to the variations in the quantum efficiency of a pixel, and combined with the undersampling of the PRF, it is most severe in channel 1. The correction can be as much as 4% peak to peak. The effect is graphically shown in Figure 5.1 where the normalized measured flux density (y-axis) is plotted against the distance of the source centroid from the center of a pixel. The correction for channel 1 can be calculated from

$$Correction = 1 + 0.0535 \times [\frac{1}{\sqrt{2\pi}} - p] \tag{5.14}$$

where p is the pixel phase (p $= \sqrt{(x - x_0)^2 + (y - y_0)^2}$), where x,y is the centroid of the point source and x_0 and y_0 are the integer pixel numbers containing the source centroid. The correction was derived from photometry of a sample of stars, each star observed at many positions on the array. The "ratio" on the vertical axis in Figure 5.1 is the ratio of the measured flux density to the mean value for the star. To correct the flux of a point source, calculate the correction from Equation 5.14 and divide the source flux by that correction. Thus, the flux of sources

well-centered in a pixel will be reduced by 2.1%. Pixel phase corrections for other channels, if necessary, and after they have been more accurately determined than currently, will be given in future Data Handbook versions.
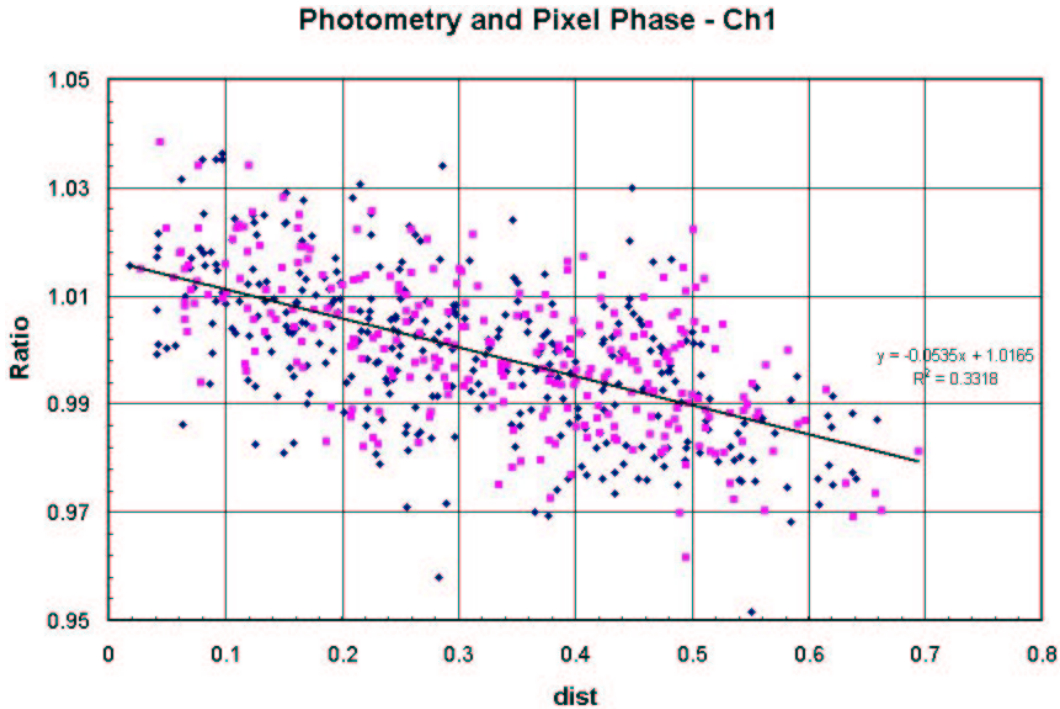

Figure 5.1: Dependence of point source photometry on the distance of the centroid of a point source from the nearest pixel center in channel 1. The ratio on the vertical axis is the measured flux density to the mean value for the star, and the quantity on the horizontal axis is the fractional distance of the centroid from the nearest pixel center.

5.5 Point-Spread Function

IRAC provides diffraction-limited imaging internally; the IRAC optics specifications limit the wavefront errors to $< \lambda/20$ in each channel. Image quality is limited primarily by the Spitzer telescope. The majority of the IRAC wavefront error is a lateral chromatic aberration that is most severe at the corners of the IRAC field. The aberration is due to the difficulty of producing an achromatic design with a doublet lens over the large bandpasses being used. The main effect is that the PSF and distortion may be slightly color dependent, which may be detectable for sources with extreme color variations across the IRAC bands.

In-flight IRAC point-spread functions (PSFs) are currently being generated from observations of a bright star observed specifically to reconstruct the PSF and will be placed on the IRAC website. The PSF currently on the IRAC web site was reconstructed from early calibration observations in which a bright star was observed multiple times with sub-pixel offsets. The data were combined onto a grid with 0.3 arcsec pixels, 1/4 the size of the IRAC detector pixels, using a 'drizzle' algorithm and a 0.6 arcsec drop size.

5.5.1 Aperture Corrections for Point Sources

The flux of point sources in IRAC images can be measured using optimal PSF fitting, aperture photometry, or other methods. In order to calibrate these methods, one uses an aperture
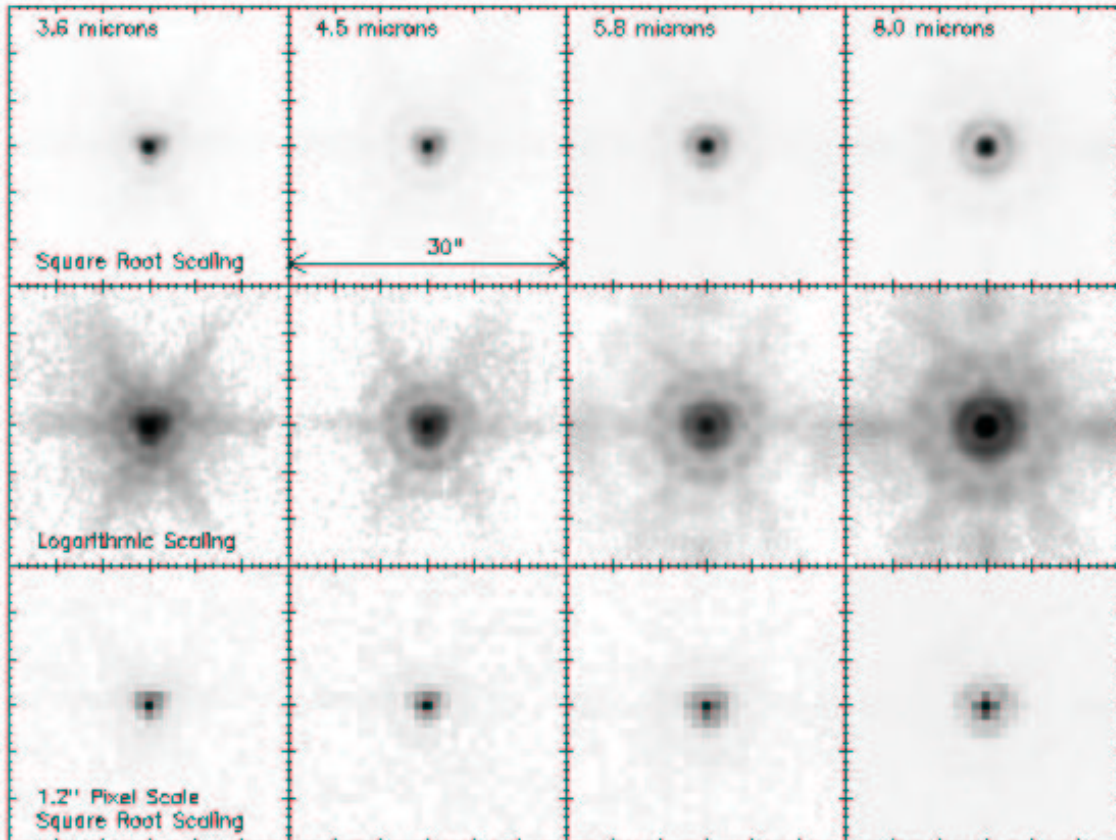
Figure 5.2: The IRAC PSF with both a square root and logarithmic scaling, emphasizing the structure in the core and wings of the PSF, respectively. We also show the PSF as it appears at the IRAC pixel scale of $1\rlap{.}''22$.

correction. The IRAC data are calibrated using aperture photometry on a set of stars. The calibration aperture has a 10 native pixel radius (12 arcsec) in all 4 channels. For crowded fields, a much smaller on-source aperture is needed. And in the presence of extended emission, a small off-source annulus is normally used. The calibration aperture does not capture all of the light from the calibration sources, so the extended emission appears too bright in the data products we delivered. See the more detailed discussion under 5.8. Similarly, observers will often use smaller apertures and will want to correct their photometry to match the absolute calibration.

Users should note that the spatial extent of the PSF in channels 3 and 4 is much larger than the subarray area. In other words, a large amount of the total power in the PSF is scattered onto arcminute size scales. As a result, special care needs to be taken when measuring fluxes in these channels, since accurate measurement of the "background" is difficult. Proper application of aperture corrections is very important.

For photometry using different aperture sizes, the aperture correction can be estimated with Table 5.7. All distances in this table are in native pixels (1″.22). The aperture corrections as written will INCREASE the flux measured by the listed method, i.e. your measured brightness should be MULTIPLIED by the aperture corrections in the table. The third decimal place in these numbers is included only to illustrate the trends; the accuracy of these corrections is presently $\sim 1\%$–2%. The aperture corrections in Table 5.7 are averages of the values derived from PSFs measured using stars at 23 different positions on the array. Standard deviations (including measurement errors and true variations across the array) are less than 0.5% for all entries except the smallest aperture, where they are still less than 1%. The extended source (infinite) corrections in Table 5.7 come from Reach et al. (2005).

Table 5.7: IRAC aperture corrections

radius on source	background annulus	Aperture correction			
		3.6 μm	4.5 μm	5.8 μm	8.0 μm
infinite	N/A	0.944	0.937	0.772	0.737
10	10-20	1.000	1.000	1.000	1.000
5	10-20	1.049	1.050	1.058	1.068
5	5-10	1.061	1.064	1.067	1.089
3	10-20	1.112	1.113	1.125	1.218
3	3-7	1.124	1.127	1.143	1.234
2	10-20	1.205	1.221	1.363	1.571
2	2-6	1.213	1.234	1.379	1.584

The IRAC PSF is available on the SSC website so that any custom flux-measurement algorithm can be compared to our calibration aperture to derive the appropriate aperture corrections. As the IRAC PSF is slightly undersampled (particularly at 3.6 and 4.5 microns), observers should derive their own aperture corrections when performing photometry on mosaics.

5.5.2 PSF Removal

The IRAC team is currently performing detailed experiments with building and removing point-spread functions from images. Until a versatile, robust PSF is finished and placed on the website, the IRAC team's best recommendation is making your own PSF from your data using post-BCD software available on the web. A second option is STinyTim, also available on our website, but it uses pre-flight numbers and is currently being updated. A third option is a PSF determined from calibration stars available on our website. The data from which the PSFs have been constructed are in the public archive.

5.6 Astrometry and Pixel Scales

5.6.1 Optical Distortion

IRAC has a $\approx 1\%$ distortion in all channels due principally to being offset from the optical axis of Spitzer. Coefficients to remove this distortion to $\sim 0''.1$ accuracy are stored in the BCD header.

Due to the telescope and camera optics, the pixel sizes of the IRAC detectors vary across each FOV in a manner that is referred to as the optical distortion. There are also more subtle effects such as a skew (the two axes are not exactly perpendicular) and difference in the pixel scales between the two axes. All of these effects are accounted for in the WCS encoded in the BCD.

Optical distortion in each of the IRAC FOVs is described in the headers using a standard method described by Shupe in

http://ssc.spitzer.caltech.edu/postbcd/doc/shupeADASS.pdf.

This method places the center of the distortion at the center of each detector array, in particular at CRPIX1 and CRPIX2. The linear terms and any skew are represented in the CD matrix header keywords (CD1_1, CD1_2, CD2_1, and CD2_2), while the distortion keywords provide the second and higher order terms. There is also a separate set of keywords representing the "reverse" distortion parameters.

This form of the optical distortion that is encoded in the BCD is read properly by several "standard" tools available to the general astronomical community: (1) the SSC mosaicker (MOPEX) in the post-BCD tools, (2) WCSTOOLS by Doug Mink (SAO) and (3) DS9 (except for grid overlays).

The optical distortion is fit independently for each IRAC detector. Originally a second-order fit was used, but an improved fit to third order was derived from the GOODS data by S. Casernato, and the coefficients from this fit have been placed in the BCD headers of data processed through the S12 and later pipelines.

5.6.2 Pixel Solid Angles

IRAC is located slightly off the optical axis of the *Spitzer* telescope. As a result there is a slight amount of distortion in the field of view. The pixels do not all subtend the same projected solid angle on the sky. The variation in pixel solid angle size is roughly 1.5%.

It is important to be aware of this because the BCD images are calibrated in units of surface brightness (specifically MJy/sr). In order to properly measure fluxes, one must take into account the size of the pixels. Failure to do so could induce photometric errors at the 1% level, depending on location on the array. Unfortunately, only the newest photometry software can read the new FITS-standard WCS distortion keywords written in the BCD headers and properly account for the sizes of the pixels. Instead, virtually all software assumes that the pixels are in units of flux per pixel, which is equivalent to units of (flux density)/sr if the pixels are the same size.

The simplest solution to this problem is to reproject the images onto an equal (or nearly so) projection system using suitable software that can understand the distortion keywords in the WCS header (e.g., MOPEX). MOPEX also has the significant advantage that it understands how to properly handle surface brightness images during coaddition. After processing, the pixels will all subtend the same solid angle, and hence any standard photometry software can produce the correct result.

However, some observers might prefer alternative approaches, in particular if they wish to measure photometry directly from the BCD images. Therefore, we supply maps of the pixel size

in the IRAC section of the SSC web pages that can be used to correct the pixel solid angles in BCD images if any measurements are being directly made on them.

5.7 Point Source Photometry

Photometry using IRAC data is no different from that with any other high-quality astronomical data. Both aperture photometry and PSF-fitting have been tried at the SSC and work successfully. Aperture photometry is most commonly used, so we will discuss it briefly. The radius of the on-source aperture should be chosen in such a way that it includes as much of the flux from the star (thus, greater than 2 arcseconds) as possible, but it should be small enough that a nearby background annulus can be used to accurately subtract unrelated diffuse emission, and that other point sources are not contributing to the aperture. For calibration stars, an annulus of 12 arcseconds is used; such a wide aperture will often not be possible for crowded fields. The dominant background in regions of low interstellar medium (ISSA 100 μm brightness less than 10 MJy/sr) is zodiacal light, which is very smooth. In regions of significant interstellar emission, it is important to use a small aperture, especially in IRAC channels 3 and 4, where the interstellar PAH bands have highly-structured emission. For example, an aperture in a star-forming region might have a radius of 3 arcseconds with a background annulus from 3 arcseconds to 7 arcseconds. The flux of a source can then be calculated in the standard way, taking the average over the background annulus, subtracting from the pixels in the on-source region, then summing over the on-source region. Aperture corrections from Table 5.7, ranging from 12% to 23%, should then be applied. The fluxes can then be converted to magnitudes, if desired, using the zero-points in Table 5.1.

5.8 Extended Source Photometry

At the current time the measurement of extended sources with IRAC remains highly problematic, particularly at 5.8 and 8 microns. The photometry of IRAC is tied to point sources (calibration stars) measured within a fiducial aperture with a radius of 12 arcseconds. Because this method is precisely characterizable and well-behaved, it results in a highly accurate calibration for point sources. However, transferring this calibration to extended sources is much more difficult, due in large part to a significant scattered component of the PSF. A significant fraction of the flux for a point source observed in IRAC channels 3 and 4 is scattered on size-scales similar to the array size, and this scattered flux mimics an apparent zero level offset.

If one bracketing case is that of a large (12 arcsecond) aperture surrounding a point source, then the other bracketing case is that of an aperture (region of interest) that is very small relative to a uniform, very extended surface brightness distribution. An example might be the surface brightness of a region a few arcseconds across in a large nebula or an extended emission cloud.

The IRAC images are in surface-brightness units (1 MJy/sr = 10^{-17} erg s^{-1} cm^{-2} Hz^{-1} sr^{-1}). The flux of an extended object is the integral of the surface brightness over the solid angle of the object. The value of a pixel in an IRAC image is not the total sky brightness, however, because the celestial "dark" subtracted from each image has some true background in it. An estimate of the total sky brightness that was seen by a pixel is the pixel value *plus* the zodiacal background from the celestial "dark" that was subtracted from a given BCD. The (theoretically) estimated zodiacal light brightness during the sky dark observation is in the BCD header keyword SKYDRKZB. Again, however, it must be emphasized that IRAC is not an absolute photometer. While it is possible using the above keyword to recover something similar to the absolute sky surface brightness, this brightness estimate is still limited by the accuracy of the underlying model of the zodiacal emission.

However, in practice most extended source photometry will usually be performed with respect to a background region within the image (for example, galaxies, nebulae, etc.) and one does not attempt to measure the absolute sky brightness on large scales (like the zodiacal cloud). While in principle the latter can be performed, it requires highly specialized observation and analysis techniques. The median value of the pixels located in user-selected background regions is generally a reasonable estimator of the background.

In this case, the flux contained in the aperture is an overestimate of the true flux. The reason is that while photons are scattered out of the aperture due to the PSF, photons are also scattered *into* the aperture from the emission regions from outside. As a thought experiment, one can imagine a single point source inside an aperture, which is easy to measure. But if four point sources are placed around it just outside the measurement aperture, each of them scatters light into the aperture and one then measures too much flux. For the extended source case, we can imagine the same experiment taken to the limit where all regions have emitters in them. Fortunately, this case is well-posed. The calculated flux must then be scaled by the ratio of the extended and point source throughputs. The scaling factors (f_p/f_{ex}) to be used are given in Table 5.7 (the infinite aperture case). Note that these are not really "throughputs," in the sense that they have anything to do with the number of photons reaching the detector. It is more accurate to think of them as a special type of aperture correction. Note also that the "extended" value applies specifically to the limiting case of measuring a very extended object in a small aperture. More specifically, the values in Table 5.7 are for a very extended, red source like the Zodiacal light.

To further complicate matters, there is a spectral dependance of the internal scattering in the Si:As arrays. The exact wavelength dependence is still being quantified through ground tests of spare arrays. The amount of scattering is most likely different for different wavelengths. Experiments indicate it is less for blue sources (stars, etc.) than red sources (zodiacal emission, HII regions). In general, the correction that needs to be applied for the photometry of an extended source is a function of source geometry and spectral energy distribution. Observers are cautioned to be careful in performing photometry on extended sources such as resolved galaxies.

In the case of extremely large apertures relative to the target size, the problem degenerates into something resembling the point source measurement. However, if the aperture is larger than 12 arcseconds, an aperture correction which decreases the measured flux may be needed because the large aperture captures more of the scattered emission than the fiducial aperture. Derivation of this aperture correction can either be made using the SSC-released PSFs, or better still could be measured directly from the image if suitable point sources exist.

Another commonly encountered problem is that of measuring the total flux of small extended objects. For example, the background galaxies seen in all IRAC images are often slightly extended on size-scales of a few arcseconds. Point source photometry of such objects will obviously underestimate their fluxes. One methodology for handling such sources was developed by the SWIRE project; readers are referred to the data release document for SWIRE, available from the SSC website under Legacy projects (go to http://ssc.spitzer.caltech.edu/legacy/original.html and click on the link to Enhanced Data Products for SWIRE, then 20050603_enhanced_v1/, then Documents). Detailed analysis by SWIRE has indicated that Kron fluxes, with no aperture corrections applied, provide measurements of small extended sources that agree closely with hand-measured fluxes, an unsurprising result given the underlying definition of the Kron flux. Kron fluxes are provided as one of several flux measures in the popular "SExtractor" software. Note that it is important to determine that an object actually is extended before using the Kron flux, as it is ill-defined otherwise. This may be determined by using the stellarity and isophotal area as defined by the SExtractor software. Selecting limits on these parameters based on their breakdown as a function of signal-to-noise ratio generally will mimic SExtractor's own "auto"

function.

The very hardest case is photometry of objects with sizes that are on arcminute scales, within apertures of similar size. Examples might be typical elliptical galaxy observations. In this case the aperture correction is related both to the aperture size and the underlying surface brightness distribution of the target. The correction will lie intermediate to the cases given in Table 5.7. At the current time users should consider several approaches, including a careful examination of the PSF and convolutions of it with their data. The IRAC-IST is actively pursuing this issue, which remains one of the largest outstanding calibration problems of IRAC.

Chapter 6

Pointing performance

Pointing is controlled by *Spitzer*'s Pointing Control System (PCS). This uses a combination of a star tracker and gyros to locate and control the attitude of the spacecraft. Absolute pointing is controlled by the star tracker, through a filter (known as the "observer") which smooths the raw star tracker output. Slews under control of the observer take ~ 10 seconds to settle, so only the initial slew and cluster slews in celestial coordinates are carried out using the observer. Once the observatory has taken the initial frame at the starting position, attitude control is handed over to the gyros. Mapping and dithering slews are made under gyro control with a shorter (≈ 5 sec) settle time. The price to be paid for the shorter settle time is that the spacecraft attitude will slowly drift with respect to the observer attitude, at a rate ~ 1 mas/sec. For long integrations (100 sec frametime), attitude control is returned to the observer after 80 sec to halt the drift.

In addition, attitude resets are performed regularly (currently every 30 min) to return the spacecraft attitude to the observer attitude. The system is designed to ensure that any motion to return the spacecraft attitude to that of the observer does not take place during an IRAC integration, to avoid smearing the images. Throughout the first 18 months of the mission the PCS system and the corresponding parts of the IRAC AOT were being adjusted for optimal performance, but below is a guide to the astrometric accuracy and image quality that might be expected from a typical observation.

6.1 Pointing Accuracy

Slews under observer control settle to the accuracy to which the star tracker to IRAC offset is known, about $0\rlap{.}''5$. Offsets between dither/mapping moves are accurate to $0\rlap{.}''1$ relative to the commanded move for small moves ($\sim 10''$), and for large moves (~ 0.5 deg) the accuracy is $\sim 0\rlap{.}''5$ (though this has been improved as of Spring 2005, and should now be only $\sim 0\rlap{.}''2$). An additional pointing error comes from the gyro drift which can accumulate over the 30 minute period between attitude resets. This error is typically $\approx 2''$ for a "worst case" frame just before a reset.

The pointing of each frame as reported in the header keywords CRVAL1 and CRVAL2 is an average of the observer attitude during the frame, and is typically accurate to $\approx 0\rlap{.}''5$ (though may be slightly worse for short frames where the observer has not fully settled). Other header keywords related to pointing include RA_RQST and DEC_RQST, the requested RA and Dec of the frame, and PTGDIFF, the difference between the requested and actual pointing. USEDBPHF should be T for all frames, if not, then pointing transfer will have failed for the frame and the user should contact the Helpdesk.

The post-BCD pointing refinement module is run by default in the post-BCD pipeline to refine the pointing to 2MASS accuracy ($\approx 0\rlap{.}''15$), and will be successful if there is a sufficient number of 2MASS stars in the data. The module operates by matching common stars between

frames ("relative refinement") and a fiducial set of stars from 2MASS ("absolute refinement"). The (RA, Dec) position and twist of each frame is then adjusted until a global minimum in the residuals is found. Application of this to the Extragalactic IRAC First Look Survey data results in a mean position error for high signal-to-noise stars with respect to 2MASS positions of $0\rlap{.}{''}25$.

The pointing refinement module writes several new keywords to the header. RFNDFLAG is true if pointing refinement was run and produced a refined solution. The refined position is given by keywords RARFND and DECRFND, and rotation by CT2RFND. A new version of the CD matrix, given by keywords CD11RFND, etc., is also written to reflect the new rotation angle (note that the pixel scale and distortion are not changed by pointing refinement). If pointing refinement fails, then the header keyword RFNDFLAG will be false and RARFND, DECRFND and CT2RFND will be set to CRVAL1, CRVAL2 and CROTA2, respectively. Note that the refined solution may be poor if the number of astrometry stars in the frame, NASTROM, is low (i.e. 0, or only a few stars). The refined pointing keywords are used by the post-BCD software if USE_REFINED_POINTING = 1 in the namelists. To use the refined pointing with other software, copy the non-standard keywords to their FITS standard equivalents, e.g. RARFND to CRVAL1, CD11RFND to CD1_1 etc. At the time of writing, pointing refinement works well on most channel 1 and 2 data, though short frames in fields near the Galactic poles in channels 3 and 4 will frequently have too few stars for a good solution.

Data processed with the S13 software version and later may have the "superboresight" correction applied (see Section 3.3.2). Users wishing to make use of the superboresight solutions need to set USE_REFINED_POINTING = 0 in MOPEX, as the superboresight pointing is contained in the standard CRVAL1,CRVAL2 and CD matrix keywords. Data which has had this correction applied will also have the ORIG_RA and ORIG_DEC keywords present which contain the initial (uncorrected) pointing estimate.

6.2   Jitter and Drift

Jitter is typically $0\rlap{.}{''}1$. It has been measured on timescales ~ 0.04 seconds to 5 minutes. In addition to high frequency jitter, there are modulations $\sim 0\rlap{.}{''}1$ on timescales of 200–400 seconds. We do not expect these to noticeably affect the IRAC PSF. Gyro drift occurs for the first 80 seconds of IRAC integrations, but again this should result in only $\approx 0.1''$ of motion. Some amount ($\leq 0\rlap{.}{''}4$) of image smearing is expected in short frames due to settling motions. Other instances of pointing glitches occur when one of the four reaction wheels goes through zero velocity. To reduce stiction when the speed actually hits zero, the wheels are given a small "bump" in torque at this point, which has been seen to result in small ($\approx 0\rlap{.}{''}5$), short duration (~ 10 seconds) pointing glitch. On average, only about one crossing per hour occurs, and they are thought to mostly happen during slews, so they are not expected to affect many IRAC images.

One manifestation of settling, jitter and drift during integrations is that the pointing of HDR short, medium (for 100s HDR) and long frames are slightly different (the same is also true of repeats taken in the same position). These differences are usually ~ 0.1 arcseconds, so they should not be a problem for most observers, but they are large enough to show up as residuals in difference images.

Chapter 7

IRAC Post BCD Processing with the SSC Software

All processing of IRAC data beyond the individual, calibrated, 256×256 images produced by the BCD science pipeline is called "post-BCD." This includes combining all images from an AOR into a mosaic, detecting sources, and any cosmetic corrections to the images (individual or mosaic) that are not based on understood instrumental artifacts or detector physics. Two important post-BCD processes are performed routinely by the pipeline and generate results that are placed in the science archive. These include Pointing Refinement, wherein a set of point sources are identified in the images and are astrometrically matched to the 2MASS catalog (see Chapter 6, and Mosaicking, wherein the individual images in an AOR are robustly combined into celestial coordinate mosaics for each IRAC waveband. The post-BCD processing can (and should) be performed in different ways for different observing strategies and scientific goals. The post-BCD pipeline processing was performed with a specific, conservative, set of parameters. Observers and archival researchers will very likely need to do post-BCD processing on their own. Most common will be generating mosaics from data in multiple AORs. The post-BCD processing software is available from the SSC website, together with extensive documentation at http://ssc.spitzer.caltech.edu/postbcd/. Here we discuss some IRAC-specific issues.

The post-BCD software consists of a series of modules linked by Perl wrapper scripts and controlled by namelists. Namelists need to be placed in a subdirectory called cdf and have filenames ending .nl. The namelist controls which modules are called, contain the names of the input file lists and output directories, and detailed parameter sets for each module. Input file lists should not have any blank lines, otherwise the programs will look for non-existent files. An example namelist for the mosaicker is given in Appendix C. The following subsections deal with each part of the post-BCD pipeline in turn, starting with pointing refinement, then overlap correction, mosaicking and finally point source extraction.

7.1 Pointing Refinement

Pointing refinement corrects the pointing of each frame to the 2MASS sky, and is carried out by default in the pipeline. It may not always be successful, however, particularly in short frames at high Galactic latitudes in channels 3 and 4, where there are few 2MASS stars per image. In these cases it may be possible to use multi-band refinement to tie the astrometry of channels 3 and 4 to that of channels 1 and 2. *Note: rerunning pointing refinement will overwrite any previous solutions in the BCD header. Users should therefore make copies of BCDs before rerunning pointing refinement if they wish to retain the old corrections.*

7.2 Overlap Correction

The post-BCD software contains an overlap correction module which matches the background levels of overlapping frames in a mosaic. This is not currently run in the automatic, online post-BCD pipeline. But generating new mosaics by running `mopex` with overlap correction turned on can remove the "patchiness" often seen in mosaics due to bias fluctuations in the array (due both to the first-frame effect and bright source effects).

7.3 Mosaicking of IRAC data

7.3.1 Creating a Common Fiducial Frame

As a first step in creating a mosaic, run the mosaicker with all the modules turned off except for the fiducial frame module (i.e. `run_fiducial_image_frame = 1` in the namelist), and include all the files you wish to mosaic from all four channels in the input list. This will generate the boundaries of the mosaic and will allow all channels to be mosaicked onto the exact same grid. Rename this file to, e.g., `FIF_all.tbl` to prevent it being accidently overwritten. Then, when running the mosaicker, set the fiducial frame to the file created by the fiducial frame module in this initial run using

    FIF_FILE_NAME = (path)/FIF_all.tbl,

and turn off the fiducial frame module (i.e.

    set run_fiducial_image_frame = 0

in the namelist). The use of the common fiducial frame will ensure that the mosaics from all four channels will be accurately co-aligned. The pixel scale is controlled by the `MOSAIC_PIXEL_RATIO_X/Y` parameters. Set `Edge_Padding = 100` to get a good border around the image. You can also specify CROTA2 for the output mosaic if you wish, or set `CROTA2 = 'A'` to get the smallest possible mosaic. The pixel size in the mosaic produced by the pipeline is exactly 1.2 arcsec × 1.2 arcsec (CDELT1, CDELT2 = ±0.000333333).

7.3.2 Bitmask Values

We currently recommend the following bitmasks:

```
DCE_Status_Mask_Fatal_BitPattern = 32544 # (bits 5, and 8-14 inclusive)
PMask_Fatal_BitPattern =  32640        # (bits 7-14 inclusive)
RMask_Fatal_BitPattern = 7
```

These values address bits which we believe are set correctly, or which might be set in future masks. Set the mask files to use with the `PMASK_FILE_NAME` (for the pmask) and `DCE_STATUS_MASK_LIST` (for the list of dmask files) keywords in the namelist.

7.3.3 Outlier Rejection

The mosaicker has three outlier rejection strategies. Single-frame outlier rejection, dual-outlier rejection and multi-frame outlier rejection. For IRAC, the most useful are the dual outlier and multi-frame rejection. Be sure to set `THRESH_OPTION = 2` in the namelist in the multi-frame &MOSAICOUTLIERIN section. Setting the thresholds too low in the outlier modules can

result in unwanted rejection of pixels in the cores of real objects. Users of these modules should carefully check the coverage maps produced by the mosaicker to ensure that the centers of real objects are not being masked out.

The outlier rejection modules set bits in the rmasks. Bit 0 is set by the single-frame outlier rejection, bit 1 by the temporal (multi-frame) rejection and bit 2 by the dual outlier rejection. Which rmask bits are used by the mosaicker is controlled by the RMask_Fatal_BitPattern. The mosaicker takes the rmasks and makes an rmask mosaic. The switches USE_OUTLIER_FOR_RMASK and USE_DUAL_OUTLIER_FOR_RMASK control which of the outlier detection modules are used. An important parameter is RM_THRESH in &MOSAICRMASKIN. When the rmask is made it is projected back onto the input images to determine which pixels will be masked in the final mosaic. An rmask pixel is divided amongst the overlapping pixels in the input image. Input image pixels with projected values of the rmask mosaic above RM_THRESH have the multi-frame outlier bit (1) set in their rmask. For channels 1 and 2 a fairly high value (e.g. 0.5) can be used. The more diffuse radhits in channels 3 and 4 can be more effectively rejected by setting RM_THRESH to a lower value, e.g. 0.05, which has the effect of growing the rmask and thus rejects the more diffuse edges of the radhit.

All the outlier rejection modules require uncertainty images. The BCD uncertainty images are adequate for this purpose, and the list of them should be specified with the `SIGMALIST_FILE_NAME` keyword in the namelist. To use them set

`have_uncertainties = 1` and

`compute_uncertainties_internally= 0.`

To compute your own uncertainty images, set

`compute_uncertainties_internally = 1,`

`have_uncertainties = 0,`

and set the appropriate values in the `&SNESTIMATORIN` section of the namelist (see the mosaicker documentation for more details).

Outlier rejection creates another set of masks, the rmasks. These indicate the pixels flagged by outlier rejection, and which are used by the mosaicker. A mosaic of the rmasks can be provided by setting `run_mosaic_rmask = 1`.

As a check on the outlier rejection, it is often helpful to examine the coverage maps output by the mosaicker. If the outlier rejection has been over-zealous there will be reductions in coverage at the positions of real sources in the mosaic. Blinking the mosaics and coverage maps in, e.g., DS9 can thus be very helpful for determining whether the outlier rejection is set up correctly to reject only genuine outliers.

7.3.4 Mosaicker Output Files

The output directory structure after running the mosaicker looks like:

Coadd	DualOutlier	Interp	ReInterp
Combine	FIF.tbl	Medfilter	Rmask
Detect	fif.tbl	mosaic_outlier.nl	Sigma
Dmask	header_list.tbl	Outlier	

The directory "Combine" contains the mosaic, mosaic.fits, a coverage map, mosaic_cov.fits and an uncertainty map mosaic_unc.fits.

7.3.5 To Drizzle or Not to Drizzle?

The mosaicker has three interpolation options, set by the `INTERP_METHOD` keyword. The default is a linear interpolation. Drizzling is available as an option, as is a grid interpolation (useful for creating quick mosaics if the PSF quality is not important). Our experience with the drizzle option suggests that it is effective when used on datasets with many dithers per sky position, and it can reduce the point-response function (PRF) width by 10%-20%, though at the expense of an unevenly-weighted image. The coverage map produced by the mosaicker can be used to investigate the pixel-to-pixel variation in the coverage of the drizzled image.

7.3.6 Mosaicking Moving Targets

Although Spitzer does track moving targets to a sub-pixel accuracy, the BCD pipeline only produces mosaics of IRAC data in fixed celestial coordinates. The user may opt to generate his or her own mosaic in a moving coordinate reference frame by setting the appropriate flags in MOPEX. The individual BCD's should be cosmetic and overlap-corrected first and then the mosaicker should be run with the flag MOVING_OBJECT_MOSAIC=1 set, using outlier rejection. Stars in the frames may be removed by outlier rejection, and the resultant composite of a moving target will be produced.

7.4 Source Extraction

The SSC source extractor, APEX, may be used to fit point sources in IRAC data. This software can be run in two modes, a single frame mode (`apex_1frame`) which can be run on an individual BCD or a mosaic, and a multi-frame mode which uses the mosaic to detect sources, but the individual BCD's to measure their fluxes.

7.4.1 Noise Estimation

The accuracy of the fluxes from APEX is very sensitive to the noise estimates, as these affect the fitted background value. For crowded fields, it is essential to include an estimate of the confusion noise (currently not included in the BCD uncertainty image). This can be estimated by measuring the difference between the actual image RMS and the estimated RMS in the uncertainty image, and then either adding it to the uncertainty image, or using it as the confusion noise value when generating uncertainty images with the post-BCD software (see above).

7.4.2 PRF Estimation

The PRFs released with the most recent version of MOPEX should be fairly good matches to the data and a significant improvement on the previous versions. Nevertheless, for high accuracy applications, observers may wish to build their own PRFs using the "prf_estimate" tool in the MOPEX package. The parameter PRF_file_name is used to specify the PRF in APEX. The PRF resample factors $PRF\_RESAMPLE\_X/Y$ need to be set in two places in the namelist (in &POINTSOURCEPROB and &SOURCESTIMATE), and must be set identically in both.

7.4.3 Background Estimation

The namelist parameter `Background_Fit` controls the type of background used for PRF fitting. If you give `Background_Fit = 0`, a median background is computed for the whole frame. A more accurate background estimate for PRF fitting, local to the source, can be generated by setting `Background_Fit = 1`. (Note that the aperture fluxes reported by APEX are always made using the median background, and hence may be inaccurate for faint sources.

7.4.4 Source Extraction

Source detection and extraction are controlled by the parameters `Detection_max/min_area` and `detection_threshold`.

APEX will frequently try to split bright sources into several components. This tendency can be controlled by setting the `Max_Number_PS` parameter in &SOURCESTIMATE to a low number (2–3).

Two files are output, extract_raw.tbl contains all detections, and extract.tbl which is a subset of extract_raw.tbl containing the objects and fields which are selected by `select_conditions` and `select_columns`.

7.4.5 Outlier Rejection

By default, APEX will not perform outlier rejection. This can be gotten around by running the mosaicker with outlier rejection turned on and keeping the intermediate products (`delete_intermediate_files = 0`).

Chapter 8

References

Abraham, P., Leinert, C., & Lemke, D. 1997 SEARCH FOR BRIGHTNESS FLUCTUATIONS IN THE ZODIACAL LIGHT AT 25 μM WITH ISO, *Astronomy & Astrophysics*, **328**, 702.

Beichman, C. *et al.* 1988. *Infrared Astronomical Satellite (IRAS): Catalogs and Atlases, Volume 1. Explanatory Supplement* (NASA RP-1190).

Blommaert, J., et al. 2003. *The ISO Handbook. Volume II: CAM – The ISO Camera* ESA Special Publication SP-1262 (ESA: Noordwijk).

Cohen, M., Megeath, S. T., Hammersley, P. L., Martín-Luis, F., & Stauffer, J. 2003. SPECTRAL IRRADIANCE CALIBRATION FOR THE INFRARED XIII: SUPERTEMPLATES AND ON-ORBIT CALIBRATORS FOR THE SIRTF INFRARED ARRAY CAMERA *Astronomical Journal*, **125**, 2645.

Fazio, G. G., *et al.* 2004, THE INFRARED ARRAY CAMERA (IRAC) FOR THE SPITZER SPACE TELESCOPE, *Astrophysical Journal Supplement*, **154**, 10.

Fixsen, D. J., Moseley, S. H., & Arendt, R. G. 2000. CALIBRATING ARRAY DETECTORS *Astrophysical Journal Supplement*, **128**, 651.

Fruchter, A. S., & Hook, R. N. 2002. DRIZZLE: A METHOD FOR THE LINEAR RECONSTRUCTION OF UNDERSAMPLED IMAGES, *Publications of the Astronomical Society of the Pacific*, 114, 144.

Hauser, M. G., *et al.* 1998. COBE Diffuse Infrared Background Experiment (DIRBE) Explanatory Supplement, `http://lambda.gsfc.nasa.gov/product/cobe/dirbe_exsup.cfm`

Quijada, M. A., *et al.* 2004. ANGLE-OF-INCIDENCE EFFECTS IN THE SPECTRAL PERFORMANCE OF THE INFRARED ARRAY CAMERA OF THE SPITZER SPACE TELESCOPE, *SPIE Proceedings, 5487, 244.*

Reach, W. T. *et al.* 2003. THE MID-INFRARED SPECTRUM OF THE ZODIACAL AND EXOZODIACAL LIGHT, *Icarus*, **164**, 384.

Reach, W. T. *et al.* 2005. ABSOLUTE CALIBRATION OF THE INFRARED ARRAY CAMERA ON THE SPITZER SPACE TELESCOPE, *PASP*, **117**, 978.

Väisänen, P., Tollestrup, E. V., & Fazio, G. G. 2001. CONFUSION LIMIT RESULTING FROM GALAXIES: USING THE INFRARED ARRAY CAMERA ON BOARD SIRTF, *Monthly Notices of the Royal Astronomical Society*, **325**, 1241.

Chapter 9

Glossary and Acronyms

ADC - Analog-to-Digital Converter

AOR - Astronomical Observation Request - an individual observation

AOT - Astronomical Observation Template - the IRAC observing mode

BCD - Basic Calibrated Data, the data product from each DCE that has been pipeline-processed and is fully calibrated. These data should be the starting point for scientific analysis in most cases.

Campaign - Unbroken time period when an instrument is powered on. Most instrument IRAC campaigns are expected to be on the order of one week in length.

DCE - Data Collection Event, for IRAC a single 256×256 pixel image from a single detector.

FPA - Focal Plane Assembly, housing one IRAC detector

InSb - Indium Antimonide, the detector material used in the short wavelength channels of IRAC (1 and 2)

IOC - In-Orbit Checkout, the two-month testing period of the telescope following its launch

MIC - Multiple Instrument Chamber

PAO - Period of Autonomous Operation, the interval between ground contacts for uplinking commands and downlinking data, normally 12 hours

PCS - Pointing Control System

PRF - Point Response Function. The PRF is essentially the convolution of a box the size of the image pixel with the PSF

ROIC - Read-Out Integrated Circuit. It is the chip that contains the multiplexer (column and row scanners and the buses), the unit cell amplifiers, and the output amplifiers. The chip containing the detector diodes is bonded to it. It is often called simply the "mux."

SAO - Smithsonian Astrophysical Observatory, Cambridge, MA

SSC - *Spitzer* Science Center, Caltech, Pasadena, CA

Si:As - Arsenic-doped Silicon, the detector material in the IRAC long-wavelength channels (3 & 4)

SWIRE - *Spitzer* Wide Area Infrared Survey

Appendix A

Introduction to IRAF/IDL Analysis of IRAC Data

IRAC data are no different from images obtained by other direct-imaging astronomical instruments. Here we give some tips for analyzing IRAC data using two popular data reduction and analysis packages, the Image Reduction and Analysis Facility (IRAF), developed by the National Optical Astronomy Observatories (NOAO), and the Interactive Data Language (IDL), developed commercially by Research Systems Inc. IRAF is structured in terms of packages, which include collections of tasks that can perform a large variety of operations on the data. IRAF documentation and software are available at `http://iraf.noao.edu/`. IDL is a programming language that can be used to create scripts to perform from simple to very sophisticated operations on the data.

A.1 IRAF

Many useful tasks can be found in the STSDAS and TABLES packages, so we recommend that you install them: `http://www.stsci.edu/resources/software_hardware/stsdas/download`.

IRAC mosaics tend to be rather large, so you may need to change the display device within IRAF to display the full image; for example:

```
xgterm&
cl
cl> !ds9&
cl> set stdimage=imt2048
cl> imheader SPITZER_I4_0008291840_0077_0001_1_bcd.fits long+ | page
cl> display SPITZER_I4_0008291840_0077_0001_1_bcd.fits 1
cl> imexam
```

You can also change the `stdimage` parameter in your `login.cl` file. Once your image has been successfully displayed, you can examine it with the task `imexamine`, as you would normally do. To quickly look through all the data of an AOR, say in channel 1, you can do the following with IRAF and ds9 image display device:

```
cl> imstack SPITZER_I1*.fits cube.fits
```

Then outside of IRAF go into ds9, open the cube.fits file (File–>Open...) and then click on Frame–>Data cube... and 'Play'.
To get rid of NaNs in a FITS file (converting them to zeros) in IRAF:

```
cl> imcopy in.fits out.fits
cl> display out.fits 1
```

To access plane 4 in an IRAC subarray image, specify the plane: image.fits[4]. (Note that the planes are numbered starting with 0.) To split up the subarray image into 64 two-dimensional planes, use package `imutil`, task `imslice`, and `set slice_dimens = 3`.

A.2   IDL

First, be sure you have the IDL astronomy library: http://idlastro.gsfc.nasa.gov/. You can test whether it is installed by issuing the command `astro` from within IDL, which should start the extremely useful astronomical coordinate converter. To read an IRAC image, use the idlastro routine `readfits`; then display it (if you know the desired greyscale range, 3–15 in the example below) as follows:

```
IDL> image=readfits('SPITZER_I4_0008291840_0081_0000_1_bcd.fits',header)
IDL> tvscl,image>3<15
```

A useful function in the idlastro library is `sigrange`, which scales an image to minimum and maximum values based on percentile levels in the image. Many IDL mathematical routines are "troubled" by NaNs, which are used by the SSC pipeline both in BCD images (where the bad pixels are NaN:ed) and the mosaics (where pixels with no valid sky data are NaN:ed). It is easy to replace the NaNs with more benign values using `where` and `finite`. This usually gives good results, by setting the NaN:ed pixels to the median of the non-Nan pixels. The example below displays the 3×3 smoothed image:

```
IDL> lnan=where(finite(image) eq 0,cnan)
IDL> lok=where(finite(image) ne 0)
IDL> if cnan gt 0 then image(l)=median(image(lok))
IDL> tvscl,sigrange(smooth(imagesmooth,3))
```

To convolve an image with a kernel, for example to convolve a channel 4 image with the channel 1 PSF, use `convol`. The example below reads the PSF file from the SSC website, regrids it to the IRAC pixel size, convolves the two, and displays the result:

```
IDL> psf=readfits('ch1psf81pix.fits',hpsf)
IDL> psf_scaled=congrid(psf,/center,21,21)
IDL> psf_scaled=psf_scaled/total(psf_scaled)
IDL> imconv=convol(image,psf_scaled)
IDL> tvscl,sigrange(imconv>10)
```

Here is how you would create a mask of bad pixels (1s) and good (0s) from the dmask. The result is written into a file called 'newmask.fits.' This file then can be read into iraf, xstarfinder, or another package. The useful bits in the dmask are 5,6,8,9,10,11,14.

```
dmask = readfits('SPITZER_I1_8845056_31_0_1_bdmsk.fits', hdmask)
bmask = dmask * 0
bmask = byte(bmask)
ptr = where(dmask and 20320L, count) ; use bits 5,6,8,9,10,11,14
if (count gt 0) then bmask[ptr] = 1B
tvscl,bmask,256,0
writefits, 'newmask.fits', bmask, hdmask
```

To work with raw subarray data, it helps to convert the 256×256 images into cubes. Here we read in a raw subarray image, convert it into a cube, display all 64 planes, and then median them into an image.

```
m=readfits('SPITZER_I1_8845056_0_0_1_dce.fits',h)
mc=rebin(m,32,32,64)
tvscl,sigrange(mc(*,*,i)),i
medarr,mc,mm
tvscl,rebin(sigrange(mm),320,320,/sample)
```

A.3   Other Software

A.3.1   SExtractor

Here is how you could use the IRAC mask files in context of point source extraction. SExtractor can use mask images to fill out an image flags field (see section 8.2 of the SExtractor manual).

In the catalog parameter (.param) file, include the keywords IMAFLAGS_ISO and/or NIMAFLAGS_ISO. SExtractor will then look for a 8,16 or 32-bit integer fits file containing the flag map. This could be the dmask (or $2^{(dmask)}$). The flag map values within the isophotal area of a given object are then combined and stored in the long integer field IMAFLAGS_ISO, and the number of flagged pixels in NIMAFLAGS_ISO. Different ways of combining the flags can be specified with the FLAG_TYPE keyword, see the SExtractor manual for more details.

Appendix B

Sample IRAC BCD Header

```
BITPIX  =                    -32 / FOUR-BYTE SINGLE PRECISION FLOATING POINT
NAXIS   =                      2 / STANDARD FITS FORMAT
NAXIS1  =             256
NAXIS2  =             256
ORIGIN  = 'Spitzer Science Center' / Organization generating this FITS file
CREATOR = 'S13.0.1 '            / SW version used to create this FITS file
TELESCOP= 'Spitzer '            / SPITZER Space Telescope
INSTRUME= 'IRAC    '            / SPITZER Space Telescope instrument ID
COMMENT    Controlled data files (CDFs) used:
COMMENT    w_bqd_files_to_copy_to_sandbox.nl, fileID = 11092005
COMMENT    w_bqd_pointrefine.nl, fileID = 11092005
CHNLNUM =                      1 / 1 digit instrument channel number
EXPTYPE = 'flt     '            / Exposure Type
REQTYPE = 'AOR     '            / Request type (AOR, IER, or  SER)
AOT_TYPE= 'IracMap '            / Observation template type
AORLABEL= 'skyflt_zody_field_23_1' / AOR Label
FOVID   =                     81 / Field of View ID
FOVNAME = 'IRAC_FOV_bwn_3.6&5.8_and_4.5&8.0_for_mapping' / Field of View Name


          / PROPOSAL INFORMATION


OBSRVR  = 'William Reach'       / Observer Name (Last, First)
OBSRVRID=                    125 / Observer ID of Principal Investigator
PROCYCL =                      3 / Proposal Cycle
PROGID  =                   1130 / Program ID
PROTITLE= 'SIRTF IRAC Calibration Program' / Program Title
PROGCAT =                     32 / Program Category


          / TIME AND EXPOSURE INFORMATION


DATE_OBS= '2005-11-29T19:24:59.652' / Date & time at DCE start
MJD_OBS =            53703.809 / [days] MJD at DCE start (,JD-2400000.5)

UTCS_OBS=        186564299.652 / [sec] J2000 ephem. time at DCE start
SCLK_OBS=        817759742.719 / [sec] SCLK time (since 1/1/1980) at DCE start
SAMPTIME=                  0.2 / [sec] Sample integration time
FRAMTIME=                 100. / [sec] Time spent integrating (whole array)
```

```
COMMENT    Photons in Well = Flux[photons/sec/pixel] * FRAMTIME
EXPTIME =                 96.8 / [sec] Effective integration time per pixel
COMMENT    DN per pixel = Flux[photons/sec/pixel] / GAIN * EXPTIME
INTRFDLY=               12.402 / [sec] Inter Frame Delay Time
AINTBEG =             532233.13 / [Secs since IRAC turn-on] Time of integ. start
ATIMEEND=             532333.07 / [Secs since IRAC turn-on] Time of integ. end
AFOWLNUM=                   16 / Fowler number
AWAITPER=                  468 / [0.2 sec] Wait period
ANUMREPS=                    1 / Number of repeat integrations
AREADMOD=                    0 / Full (0) or subarray (1)
HDRMODE =                    F / DCE taken in High Dynamic Range mode
ABARREL =                    4 / Barrel shift
APEDSIG =                    0 / 0=Normal, 1=Pedestal, 2=Signal


         / TARGET AND POINTING INFORMATION


OBJECT  = 'zody_field_23'      / Target Name
OBJTYPE = 'TargetFixedSingle'  / Object Type
CRVAL1  =       330.632186594059 / [deg] RA at CRPIX1,CRPIX2
CRVAL2  =       -7.54912407871308 / [deg] DEC at CRPIX1,CRPIX2


RA_HMS  = '22h02m31.7s'        / [hh:mm:ss.s] CRVAL1 as sexagesimal
DEC_DMS = '-07d32m57s'         / [dd:mm:ss] CRVAL2 as sexagesimal
RADESYS = 'ICRS    '           / International Celestial Reference System
EQUINOX =                 2000. / Equinox for ICRS celestial coord. system
CD1_1   = -0.000118976974594661 / Corrected CD matrix element with Pointing Reco
CD1_2   = 0.000318407144176563 / Corrected CD matrix element with Pointing Recon


CD2_1   = 0.000318308468706923 / Corrected CD matrix element with Pointing Recon
CD2_2   = 0.000118665922217198 / Corrected CD matrix element with Pointing Recon
CTYPE1  = 'RA---TAN-SIP'       / RA---TAN with distortion in pixel space
CTYPE2  = 'DEC--TAN-SIP'       / DEC--TAN with distortion in pixel space
CRPIX1  =               128.
CRPIX2  =               128.
PXSCAL1 =     -1.22334117768332 / [arcsec/pix] Scale for axis 1 at CRPIX1,CRPIX2
PXSCAL2 =      1.22328355209902 / [arcsec/pix] Scale for axis 2 at CRPIX1,CRPIX2
CRDER1  = 0.000303165730049659 / [deg] Uncertainty in CRVAL1
CRDER2  = 0.000303120919216518 / [deg] Uncertainty in CRVAL2
UNCRTPA =  0.00167501499440497 / [deg] Uncertainty in position angle
CSDRADEC= 7.06716853809184E-06 / [deg] Costandard deviation in RA and Dec
SIGRA   =     0.049555601690536 / [arcsec] RMS dispersion of RA over DCE
SIGDEC  =    0.0446213252369505 / [arcsec] RMS dispersion of DEC over DCE
SIGPA   =     0.586026763706276 / [arcsec] RMS dispersion of PA over DCE
PA      =      69.5602738805828 / [deg] Position angle of axis 2 (E of N) (was OR
RA_RQST =      330.632188779167 / [deg] Requested RA at CRPIX1, CRPIX2
DEC_RQST=     -7.54902580312129 / [deg] Requested Dec at CRPIX1, CRPIX2
PM_RA   =                   0. / [arcsec/yr] Proper Motion in RA (J2000)
PM_DEC  =                   0. / [arcsec/yr] Proper Motion in Dec (J200)
RMS_JIT =   0.0284596885521926 / [arcsec] RMS jitter during DCE
RMS_JITY=   0.0200595620178258 / [arcsec] RMS jitter during DCE along Y
```

```
RMS_JITZ=    0.0201883095909688 / [arcsec] RMS jitter during DCE along Z
SIG_JTYZ=  -0.0149094839770305 / [arcsec] Costadard deviation of jitter in YZ
PTGDIFF =    0.35386296485865 / [arcsec] Offset btwn actual and rqsted pntng
PTGDIFFX=    -0.3287939507534 / [pixels] rqsted - actual pntng along axis 1
PTGDIFFY=   0.131174524549524 / [pixels] rqsted - actual pntng along axis 2
RA_REF  =         330.56637611 / [deg] Commanded RA (J2000) of ref. position
DEC_REF =          -7.44284361 / [deg] Commanded Dec (J2000) of ref. position
USEDBPHF=                    T / T if Boresight Pointing History File was used
BPHFNAME= 'BPHF.0817732800.01.pntg' / Boresight Pointing History Filename
FOVVERSN= 'BodyFrames_FTU_14a.xls' / FOV/BodyFrames file version used
RECONFOV= 'IRAC_Center_of_3.6umArray' / Reconstructed Field of View



                / DISTORTION KEYWORDS

A_ORDER =                    3 / polynomial order, axis 1, detector to sky
A_0_2   =          2.9656E-06 / distortion coefficient
A_0_3   =          3.7746E-09 / distortion coefficient
A_1_1   =          2.1886E-05 / distortion coefficient
A_1_2   =         -1.6847E-07 / distortion coefficient
A_2_0   =         -2.3863E-05 / distortion coefficient
A_2_1   =          -8.561E-09 / distortion coefficient
A_3_0   =         -1.4172E-07 / distortion coefficient
A_DMAX  =                1.394 / [pixel] maximum correction
B_ORDER =                    3 / polynomial order, axis 2, detector to sky
B_0_2   =            2.31E-05 / distortion coefficient
B_0_3   =         -1.6168E-07 / distortion coefficient
B_1_1   =         -2.4386E-05 / distortion coefficient
B_1_2   =         -5.7813E-09 / distortion coefficient
B_2_0   =          2.1197E-06 / distortion coefficient
B_2_1   =         -1.6583E-07 / distortion coefficient
B_3_0   =         -2.0249E-08 / distortion coefficient
B_DMAX  =                1.501 / [pixel] maximum correction
AP_ORDER=                    3 / polynomial order, axis 1, sky to detector
AP_0_1  =         -6.4275E-07 / distortion coefficient
AP_0_2  =         -2.9425E-06 / distortion coefficient
AP_0_3  =          -3.582E-09 / distortion coefficient
AP_1_0  =         -1.4897E-05 / distortion coefficient
AP_1_1  =          -2.225E-05 / distortion coefficient
AP_1_2  =          1.7195E-07 / distortion coefficient
AP_2_0  =          2.4146E-05 / distortion coefficient
AP_2_1  =           6.709E-09 / distortion coefficient
AP_3_0  =          1.4492E-07 / distortion coefficient
BP_ORDER=                    3 / polynomial order, axis 2, sky to detector
BP_0_1  =         -1.6588E-05 / distortion coefficient
BP_0_2  =         -2.3424E-05 / distortion coefficient
BP_0_3  =           1.651E-07 / distortion coefficient
BP_1_0  =         -2.6783E-06 / distortion coefficient
BP_1_1  =          2.4753E-05 / distortion coefficient
```

```
BP_1_2  =            3.8917E-09 / distortion coefficient
BP_2_0  =            -2.151E-06 / distortion coefficient
BP_2_1  =               1.7E-07 / distortion coefficient
BP_3_0  =            2.0482E-08 / distortion coefficient

            / PHOTOMETRY

BUNIT   = 'MJy/sr  '            / Units of image data
FLUXCONV=                0.1104 / Flux Conv. factor (MJy/Str per DN/sec)
GAIN    =                   3.3 / e/DN conversion
ZODY_EST=             0.1315266 / [MJy/sr] Zodiacal Background Estimate
ISM_EST =           0.003466571 / [MJy/sr] Interstellar Medium Estimate
CIB_EST =                    0. / [MJy/sr] Cosmic Infrared Background Estimate
SKYDRKZB=              0.034997 / [MJy/sr] Zodiacal Background Est of subracted s
SKYDKMED=              0.072353 / [MJy/sr] Median of Subtracted Skydark

            / GENERAL MAPPING KEYWORDS

COLUMN  =                     2 / Current column number
ROW     =                     3 / Current row number
MAPNUM  =                     1 / Current map cycle number
DITHPOS =                     1 / Current dither position

            / IRAC MAPPING KEYWORDS

READMODE= 'FULL    '            / Readout mode
COLSTEP =                  120. / [arcsec] Column step size
ROWSTEP =                  120. / [arcsec] Row step size
COLOFF  =                    0. / [arcsec] Column map center offset
ROWOFF  =                    0. / [arcsec] Row map center offset
MAPCOORD= 'Array   '            / Array or celestial coordinates
DITHSCAL= 'large   '            / Dither scale (small, medium, large)

            / INSTRUMENT TELEMETRY DATA


ASHTCON =                     2 / Shutter condition (1:closed, 2: open)
AWEASIDE=                     0 / WEA side in use (0:B, 1:A)

ACTXSTAT=                     0 / Cmded transcal status
ATXSTAT =                     0 / transcal status
ACFLSTAT=                     0 / Cmded floodcal status
AFLSTAT =                     0 / floodcal status
AVRSTUCC=                  -3.5 / [Volts] Cmded VRSTUC Bias
AVRSTBEG=            -3.5120896 / [Volts] VRSTUC Bias at start integration
AVDETC  =                 -2.75 / [Volts] Cmded VDET Bias

AVDETBEG=            -2.7578687 / [Volts] VDET Bias at start of integration
AVGG1C  =                 -3.65 / [Volts] Cmded VGG1 Bias
AVGG1BEG=            -3.2078801 / [Volts] VGG1 Bias at start of integration
```

```
AVDDUCC =                    -3 / [Volts] Cmded VDDUC Bias
AVDDUBEG=                    -3 / [Volts] VDDUC Bias at start integration
AVGGCLC =                     1 / [Volts] Cmnded VGGCL clock rail voltage
AVGGCBEG=                     1 / [Volts] VGGCL clock rail voltage
AHTRIBEG=             204.59911 / [uAmps] Heater current at start of integ
AHTRVBEG=             2.3936886 / [Volts] Heater Voltage at start integ.
AFPAT2B =             15.023121 / [Deg_K] FPA Temp sensor #2 at start integ.
AFPAT2BT=             532230.48 / [Sec] FPA Temp sensor #2 time tag
AFPAT2E =             15.023121 / [Deg_K] FPA temp sensor #2, end integ.
AFPAT2ET=             532320.26 / [Sec] FPA temp sensor #2 time tag
ACTENDT =             20.102446 / [Deg_C] C&T board thermistor
AFPECTE =              18.04231 / [Deg_C] FPE control board thermistor
AFPEATE =             21.595371 / [Deg_C] FPE analog board thermistor
ASHTEMPE=             21.291198 / [Deg_C] Shutter board thermistor
ATCTEMPE=             22.449469 / [Deg_C] Temp. controller board thermistor
ACETEMPE=             20.071965 / [Deg_C] Calib. electronics board thermistor
APDTEMPE=             20.315812 / [Deg_C] PDU board thermistor
ACATMP1E=             1.3140427 / [Deg_K] CA Temp, end integration for temp1
ACATMP2E=             1.2972159 / [Deg_K] CA Temp, end integration for temp2
ACATMP3E=             1.3292052 / [Deg_K] CA Temp, end integration for temp3
ACATMP4E=             1.3260648 / [Deg_K] CA Temp, end integration for temp4
ACATMP5E=             1.3241778 / [Deg_K] CA Temp, end integration for temp5
ACATMP6E=             1.3226423 / [Deg_K] CA Temp, end integration for temp6
ACATMP7E=             1.3213983 / [Deg_K] CA Temp, end integration for temp7
ACATMP8E=             1.3143697 / [Deg_K] CA Temp, end integration for temp8

          / DATA FLOW KEYWORDS

ORIGIN0 = 'JPL_FOS '          / Site where RAW FITS file was written
CREATOR0= 'J5.3     '          / SW system that created RAW FITS
DATE    = '2005-11-30T05:30:56' / [YYYY-MM-DDThh:mm:ss UTC] file creation date
AORKEY  =              16408832 / AOR or EIR key. Astrnmy Obs Req/Instr Eng Req
DS_IDENT= 'ads/sa.spitzer#0016408832' / Data Set Identification for ADS/journals
EXPID   =                    44 / Exposure ID (0-9999)
DCENUM  =                     0 / DCE number (0-9999)
TLMGRPS =                     1 / expected number of groups
FILE_VER=                     1 / Version of the raw file made by SIS
RAWFILE = 'IRAC.1.0016408832.0044.0000.01.mipl.fits' / Raw data file name
CPT_VER = '3.1.11   '          / Channel Param Table FOS versioN
CTD_VER = '3.0.94S '          / Cmded telemetry data version

EXPDFLAG=                     T / (T/F) expedited DCE
MISS_LCT=                     0 / Total Missed Line Cnt in this FITS
MANCPKT =                     F / T if this FITS is Missing Ancillary Data
MISSDATA=                     F / T if this FITS is Missing Image Data
CHECKSUM=                     0 / MIPL computed checksum
PAONUM  =                  1558 / PAO Number
CAMPAIGN= 'IRAC007400'        / Campaign
DCEID   =              46955729 / Data-Collection-Event ID
DCEINSID=              10091777 / DCE Instance ID
```

```
DPID    =           63087225 / Data Product Instance ID
PIPENUM =                107 / Pipeline Script Number
SOS_VER =                  1 / Data-Product Version
PLVID   =                  4 / Pipeline Version ID
CALID   =                  6 / CalTrans Version ID


SDRKEPID=            1254939 / Sky Dark ensemble product ID


PMSKFBID=                995 / Pixel mask ID
LDRKFBID=                804 / Fall-back lab dark ID
LINCFBID=                357 / Fall-back Linearity correction ID
FLATFBID=                791 / Fall-back flat ID
FLXCFBID=                951 / Flux conversion ID
MBLTFBID=                696 / Muxbleed Lookup Table ID
MBCFFBID=                704 / Muxbleed Coefficients ID
LBDRKFLE= 'FUL_100s_100sf16d1r1_ch1_v1.2.0_dark.txt' / Labdark File Used



            / PROCESSING HISTORY


HISTORY job.c ver: 1.50
HISTORY TRANHEAD               v.          12.6, ran Tue Nov 29 21:30:28 2005
HISTORY CALTRANS               v.           4.0, ran Tue Nov 29 21:30:37 2005
HISTORY INSBPOSDOM             v.           1.1, ran Tue Nov 29 21:30:38 2005
HISTORY cvti2r4       v.  1.30 A50506, generated 11/29/05 at 21:30:38
HISTORY hdrupd8       v.  1.5  A50126, updated  11/29/05 at 21:30:40
HISTORY FFCORR              v. 2.000, ran Tue Nov 29 21:30:41 2005
HISTORY MUXBLEEDCORR           v.         1.600, ran Tue Nov 29 21:30:42 2005
HISTORY FOWLINEARIZE           v.      4.900000, ran Tue Nov 29 21:30:43 2005
HISTORY BGMODEL                v.           1.0, ran Tue Nov 29 21:30:44 2005
HISTORY SLREMOVE               v.           1.0, ran Tue Nov 29 21:30:44 2005
HISTORY hdrupd8       v.  1.5  A50126, updated  11/29/05 at 21:30:44
HISTORY hdrupd8       v.  1.5  A50126, updated  11/29/05 at 21:30:45
HISTORY hdrupd8       v.  1.5  A50126, updated  11/29/05 at 21:30:45
HISTORY DARKSUBNG             v. 1.000, ran Tue Nov 29 21:30:46 2005
HISTORY FLATAP               v. 1.500   Tue Nov 29 21:30:46 2005
HISTORY DNTOFLUX               v.           4.1, ran Tue Nov 29 21:30:56 2005
HISTORY CALTRANS               v.           4.0, ran Wed Nov 30 05:59:53 2005
HISTORY PTNTRAN                v.           1.3, ran Wed Nov 30 06:00:02 2005
HISTORY FPGen                  v.          1.25, ran Wed Nov 30 06:00:03 2005
HISTORY CALTRANS               v.           4.0, ran Wed Nov 30 06:18:54 2005
SOFTWARE= 'pointingrefine'    / Pointing refinement using pnt-src correlation
PTGVERSN=                6.7 / Version number of pointingrefine program
RFNDFLAG=                  T / T if pointing was actually refined
RARFND  =      330.63213269071 / [deg] Refined RA
DECRFND =     -7.54901470197056 / [deg] Refined DEC
CT2RFND =     -69.5601949004172 / [deg] Refined CROTA2
PA_RFND =      69.5601949004172 / [deg] Refined PA (= -CROTA2_refined)
ERARFND = 1.25941341256569E-06 / [deg] Error in refined RA
```

75

```
EDECRFND= 1.20728039549886E-06 / [deg] Error in refined DEC
EPA_RFND= 0.000213490721788015 / [deg] Error in refined PA or CROTA2
CSRDRFND= -2.37991046957724E-07 / [deg] Co-standard dev. in refined RA, DEC
NASTROM =                    13 / # Astrometric sources for absolute refinement
RARESID =     0.19237012846764 / [arcsec] Residual: Observed-Refined RA
DECRESID=   -0.393756273078338 / [arcsec] Residual: Observed-Refined DEC
PA_RESID=    0.284328595995476 / [arcsec] Residual: Observed-Refined PA
XC_RESID=   -0.356519564131304 / [pix] Refinement correction in CRPIX1
YC_RESID=   0.0346051945908954 / [pix] Refinement correction in CRPIX2
EXCRESID=  0.00361645437019475 / [pix] Uncertainty in corr. for CRPIX1
EYCRESID=  0.00364417164722893 / [pix] Uncertainty in corr. for CRPIX2
CSDXCYC = -0.000857499119337334 / [pix] Co-standard dev. in ref. CRPIX1, CRPIX2
CD11RFND= -0.000118977413371292 / [deg/pix] Refined CD matrix element 1_1
CD12RFND= 0.000318406980599567 / [deg/pix] Refined CD matrix element 1_2
CD21RFND= 0.000318308304701152 / [deg/pix] Refined CD matrix element 2_1
CD22RFND= 0.000118666361129849 / [deg/pix] Refined CD matrix element 2_2
END
```

Appendix C

Sample Mosaicker Namelist

Below is an example of a mosaic namelist that was used for processing the extragalactic FLS. The selection of modules to be run is listed at the top of the file. Note that the single-frame outlier rejection (`run_detect_radhit`) is turned off (set = 0) as the coverage is large enough (5 frames per sky position) that dual outlier and multi-frame rejection, which are both more reliable than the single-frame rejection, can be used. Also, `run_fiducial_image_frame` is switched off, because in this case the mosaicker was run previously with just the `run_fiducial_image_frame` turned on and created the fiducial frame file `FIF_big.tbl` (input using `FIF_FILE_NAME = FIF_big.tbl`). Points to note: (1) module parameter sets are defined by &s in front of the module name, and end with `&END`. (2) Comments are preceded by hash marks. (3) Lines containing parameter settings must end with a comma *even if the line is commented out*. (4) Spaces must appear on both sides of equal (=) signs (5) parameter and module names are case sensitive.

```
run_medfilter = 1
#detect_radhit does single frame radhit rejection
#by default donot run
run_detect_radhit = 0
have_uncertainties = 1
compute_uncertainties_internally = 0
run_fiducial_image_frame = 0
run_mosaic_interp = 1
run_detect_outlier = 1
run_mosaic_proj = 1
run_mosaic_covg = 1
run_mosaic_dual_outlier = 1
run_level = 1
run_mosaic_outlier = 1
run_mosaic_rmask = 1
run_mosaic_reinterp = 1
run_mosaic_coadder = 1
run_mosaic_combiner = 1
NICE = 0
verbose = 1
delete_intermediate_files = 1
save_namelist = 1
overwrite_dmask  = 0

USE_OUTLIER_FOR_RMASK = 1
```

```
USE_DUAL_OUTLIER_FOR_RMASK = 1
USE_REFINED_POINTING = 1
sigma_weighted_coadd = 0

MOSAIC_PIXEL_RATIO_X = 1
MOSAIC_PIXEL_RATIO_Y = 1

DMASK_DIR = Dmask_1
INTERP_DIR = Interp_1
COADDER_DIR = Coadd
COMBINER_DIR = Combine_1
REINTERP_DIR = Reinterp_1
OUTLIER_DIR = Outlier_1
DUAL_OUTLIER_DIR = DualOuliter_1
DETECT_DIR = Detect_1
RMASK_DIR = Rmask_1
MEDFILTER_DIR = Medfilter
OUTPUT_DIR = mosaic_outlier_3
FIF_FILE_NAME = FIF_big.tbl
SIGMA_DIR = Sigma_Orig

IMAGE_STACK_FILE_NAME = chan3_bcd.lis
SIGMALIST_FILE_NAME = chan3_sig.lis
PMASK_FILE_NAME = chan3_pmask.fits
DCE_STATUS_MASK_LIST =  chan3_dmask.lis
#RMASK_LIST = mosaic_outlier_1/Rmask_1/rmask_list.txt

DCE_Status_Mask_Fatal_BitPattern = 32544
PMask_Fatal_BitPattern =  32767
RMask_Fatal_BitPattern = 7

 &FIDUCIALIMAGEFRAMEIN
#   Edge_Padding is given in terms of pixels,
   Edge_Padding = 100,
# if specified CROTA2 is used for FIF without modification,
#   CROTA2 = -9.999000e+01,
   CROTA2 = A,
 &END

 &MEDFILTER
  Window_X = 21,
  Window_Y = 21,
  N_Outliers_Per_Window = 100,
&END

&DETECT_RADHIT
# Detection_Max_Area (Optional, default is 3),
  Detection_Max_Area      = 3,
# Segmentation_Threshold (Optional, default is 3.),
  Segmentation_Threshold  = 3.,
```

```
#  Radhit_Threshold (Optional, default is 6.),
  Radhit_Threshold        = 10.,
  Detect_Small_RadHits = 1,
  Detect_Streak_1 = 0,
  Detect_Streak_2 = 0,
#  Streak_Min_Area (Optional, default is 5.),
  Streak_Min_Area = 4,
#  Streak_Thickness (Optional, default is 1.5),
  Streak_Thickness = 0.7,
&END

 &SNESTIMATORIN
   Gain = 309,
   Confusion_Sigma = 8.7,
   Read_Noise = 9,
 &END

&MOSAICINTIN
# INTERP_METHOD = 1(default), 2(Drizzle),3(Grid),
   INTERP_METHOD = 1,
# used by Drizzle,
   DRIZ_FAC = 0.8,
# used by Grid,
   GRID_RATIO = 2,
&END

&MOSAICREINTIN
&END

&DETECT
# the maximum number of pixels in a cluster to be declared a detection,
# for bigger clusters the program raises the detection threshold,
# and either splits the cluster into 2 or shrinks it down to the required size,
  Detection_Max_Area         = 100,
# the minimum number of pixels in a cluster to be declared a detection,
# include all the clusters here, since we are looking for radhits,
  Detection_Min_Area         = 0,
  Detection_Threshold        = 4,
&END

&MOSAICPROJIN
&END

&MOSAICCOVGIN
   TILEMAX_X = 500,
   TILEMAX_Y = 500,
&END

&MOSAICDUALOUTLIERIN
#  MIN_OUTL_IMAGE = 1,
```

```
#   MIN_OUTL_FRAC = 0.2,
  MAX_OUTL_IMAGE = 1,
  MAX_OUTL_FRAC = 1,
TILE_XSIZ = 500,
   TILE_YSIZ = 500,
&END

&LEVEL
  Threshold_Ratio   = 0.5,
# (Optional, default is 0.500000),
&END

&MOSAICOUTLIERIN
   THRESH_OPTION = 2,
   BOTTOM_THRESHOLD = 4.,
   TOP_THRESHOLD = 4.,
   MIN_PIX_NUM = 3,
   TILE_XSIZ = 500,
   TILE_YSIZ = 500,
&END

&MOSAICRMASKIN
  RM_THRESH = 0.5,
  MIN_COVERAGE = 3,
  MAX_COVERAGE = 100,
&END


&MOSAICCOADDIN
   TILEMAX_X = 1000,
   TILEMAX_Y = 1000,
&END
```

Index

warm electronics, 2