Spitzer Space Telescope

Observer's Manual – Version 8.0

Release Date: 15 August 2007

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

ii

Issue Dates

Acknowledgements

Version 8.0 Summary of Changes

This is a summary of the major changes in version 8.0 based upon version 7.1.

The main purpose of this release is to update the information for CP-5. Myriad small changes were made throughout.

Chapter 1 – Document list/version numbers were updated.
Chapter 2 – Updated information for CP-5.
Chapter 3 – BIC for Cycle-4 changed to ~10:11:14. MIPS-warm and MIPS-cold campaigns are continued.
Chapter 4 – No changes.
Chapter 5 – Updated information relevant to CP-5.
Chapter 6 – Various updates and modifications. Added new subsections on intra-pixel sensitivity variations and high precision relative photometric monitoring.
Chapter 7 – Various updates and modifications, primarily removed references to IRSFRINGE, which is no longer applicable.
Chapter 8 – Various updates, particularly relating to the MIPS 160 μm "enhanced" photometry mode.
Chapter 9 – No updates.

Table of Contents: Major Sections

List of Tables

viii

Table of Contents: All Sections

1 Philosophy and Scope of the Spitzer Space Telescope Observer's Manual

The Spitzer Space Telescope is the fourth element in NASA's family of Great Observatories and represents an important scientific and technical component in NASA's Astronomical Search for Origins program. The spacecraft carries an 85-centimeter telescope and three cryogenically-cooled science instruments capable of performing imaging and spectroscopy from 3 to 180 μm. Spitzer was launched on a Delta 7920H from Cape Canaveral, Florida into an Earth-trailing heliocentric orbit on 25 August 2003. While the Spitzer cryogenic lifetime requirement is 2.5 years of normal operations, which was passed on 26 April 2006, current estimates indicate an expected cryogenic lifetime of about 5.5 years.

1.1 Intended Audience

This document, the Spitzer Space Telescope Observer's Manual (SOM), is the essential technical reference manual for Spitzer observers. The SOM provides information about the design, performance and operational constraints of Spitzer. It includes information on planning, editing and submitting Spitzer observations.

Many documents, including the SOM, are available on the public Spitzer Science Center (SSC) website at http://ssc.spitzer.caltech.edu/. We have updated this document with the latest information known as of August 2007.

For readers interested in data analysis, considerable additional documentation is available on the SSC website, including data handbooks and data processing cookbooks.

1.2 Document Organization

The SOM is divided into four major parts: the Introduction (Chapters 1 and 2), the observatory chapters (Chapters 3 to 5), and the instrument-specific chapters (Chapters 6 to 8). The content of a chapter on cross-calibration, in previous versions of the SOM, is now exclusively available online at http://ssc.spitzer.caltech.edu/obs/xcalib.html. The observatory section includes a technical description of Spitzer, its operational capabilities and constraints, and information about how Spitzer observations are planned and specified.

Each instrument has its own chapter, prepared by the instrument teams and SSC instrument support teams. All three chapters follow the same basic design. They include a technical description of the instrument, an overview of the instrument's performance and capabilities, a description of its observing modes, advice on how to best use them, and a discussion of the data reduction pipeline and the characteristics of the data. More extensive discussion of the data is presented in the Data Handbooks which are available for all three instruments on the SSC website.

1.3 Document Update Plans

This manual is maintained by the SSC and is intended to support the fifth Call for Proposals (CP). Proposers should note that we have updated this document with the latest information known as of August 2007; in a very few cases, references

are made here to information that is posted on the SSC website. In addition to maintaining the SOM, the SSC maintains instrument web pages at: http://ssc.spitzer.caltech.edu/. No significant updates are anticipated between the release of this document and the CP-5 deadline.

1.4 Other Relevant Documents

If more recent versions of any of these documents are available on the SSC website, please use those versions.

- Spitzer Space Telescope Observer's Manual (this document – version 8.0)
- Spitzer Space Telescope Cycle-5 Call for Proposals (version 1.0, dated 15 August 2007)
- Spitzer Space Telescope Observing Rules (version 8.0, dated 15 August 2007)
- Spot User's Guide (to go with Spot version 16, updated 15 August 2007)
- Leopard User's Guide (to go with Leopard version 7, updated 15 August 2007)
- Spitzer Space Telescope Observation Planning Cookbook (version 7.0, dated 15 September 2007)
- Many other helpful documents can be found on the SSC website at http://ssc.spitzer.caltech.edu/.

Previous years also mentioned the Spitzer Space Telescope Reserved Observations Catalog. It is easier (and recommended) at this point in the mission to use Leopard to search the set of objects observed or planned to be observed with Spitzer. (Leopard is packaged with Spot, and also has batch search capability.) However, we are still generating the plain ASCII ROC, obtainable from the SSC website; the latest version is dated 15 August 2007.

Additional relevant documents include:
- IRAC Data Handbook (version 3.0, 20 January 2006)
- IRS Data Handbook (version 3.0, dated 5 July 2007)
- MIPS Data Handbook (version 3.2.1, 6 February 2006)

2 Introduction to Observing with Spitzer

This chapter includes an overview of the Spitzer mission, a summary of the observatory's technical capabilities, and a description of the Science User Support services that are offered by the SSC. The relevant details of the mission and telescope design appear in Chapters 3-8.

A brief, high-level summary of the mission for astronomers appears in the ApJS Spitzer Special Issue, specifically the paper by Werner et al. (2004, ApJS, 154, 1) entitled "The Spitzer Space Telescope Mission." A copy of this paper is available on the SSC website.

2.1 Categories of Observing Programs

There are four categories of observing time on Spitzer: Guaranteed Time Observations (GTOs), Legacy Science, Director's Discretionary Time (DDT), and General Observations (GOs). The Spitzer GTO program results from a 1983 NASA Announcement of Opportunity and competitive selection of instrument teams and the Spitzer Science Working Group. The GTOs were allocated 20% of the available observing time for the first 2.5 years of the Spitzer science mission, and 15% thereafter. The original Spitzer Legacy Science Program, comprised of 3160 hours of observing time, consists of six projects competitively selected in November 2000 following a solicitation of proposals and a competitive peer review. The Program was motivated by a desire to enable major science observing projects early in the Spitzer mission, with the goal of creating a substantial and coherent database of archived observations that can be utilized by subsequent Spitzer researchers. Additional Legacy Science programs were awarded in Cycles-2 through 4. Five percent of the total Spitzer observing time is allocated by the SSC Director as Director's Discretionary Time (DDT), and is intended to facilitate proposals that address emerging scientific topics. The remaining observing time, typically more than 5000 hours annually, is allocated for General Observer (GO) investigations. During Cycle-4, GO programs ranged in size from 0.2 to 292 hrs. Cycle-5 is the last cryogenic cycle for Spitzer.

2.2 Mission Overview

Spitzer was launched from Cape Canaveral, Florida into an Earth-trailing heliocentric orbit on 25 August 2003. The observatory was launched with the telescope at ambient temperature; only the focal plane instruments were cooled to cryogenic temperatures. The telescope gradually cooled to ~6 K over a period of ~45 days. A telescope temperature of ~6 K is required to bring the telescope thermal background down low enough for observing at the longest Spitzer wavelengths.

Following launch, Spitzer entered a 63-day In-Orbit Checkout (IOC) phase, followed by a 35-day Science Verification (SV) phase, during which time the planned capabilities of the telescope were verified, the detailed performance characterized, and the three science instruments and their operational modes commissioned. Following completion of IOC + SV, Spitzer was commissioned for routine science operations on 1 Dec 2003.

The prime science mission will continue until the exhaustion of the onboard cryogen used to cool the telescope and science instruments. While the Spitzer cryogenic lifetime requirement is 2.5 years of normal operations, which was passed on 26 April 2006, current estimates indicate an expected cryogenic lifetime of about 5.5 years. We expect the cryogen to be depleted by about 2009 March.

The first science observations during the nominal mission included the First Look Survey (FLS), which was conducted by the SSC on behalf of the Spitzer observer community. The FLS is a ~110-hour survey using Director's Discretionary Time. The goals of the FLS were to provide a characteristic first look at the mid-infrared sky and to rapidly process the data and place it in the public domain in time to impact early Spitzer investigations (specifically, Cycle-1 planning). A detailed description of the motivation and observation planning for the FLS is available on the SSC web pages. The first science also included GTO and Legacy Science observing. About 6 months after IOC+SV was completed, Cycle-1 GO observing commenced (June 2004). The majority of the original Legacy Science observations (which are made public immediately) were completed within the first 18 months of the science mission. The science Archive opened in May 2004, and initially contained FLS data, early release observations, and any Legacy data taken and reprocessed with the most current pipelines at that time. Cycle-2 GO observations commenced in June 2005. Cycle-3 GO observations commenced in June 2006. Cycle-4 GO observations commenced in July 2007.

This Call for Proposals is being issued in August, with proposals due in November.

2.3 High Level Observatory Description: What Can Spitzer Do?

Spitzer is a 3-axis stabilized pointing and scanning observatory. The top-level observatory characteristics are summarized in Table 2.1. The observatory as a whole is presented in more detail in Chapter 3, the telescope optics in Chapter 4 and the science instruments in Chapters 6–8.

Spitzer's science payload consists of three cryogenically-cooled instruments, which together offer observational capabilities stretching from the near- to the far-infrared. Table 2.2 gives a summary of the characteristics of each instrument, *including limits due to confusion*.

The InfraRed Array Camera (IRAC) – Giovanni G. Fazio, Smithsonian Astrophysical Observatory/Harvard-Smithsonian Center for Astrophysics, PI IRAC provides images at 3.6, 4.5, 5.8 and 8.0 μm, with two adjacent 5.2′ × 5.2′ fields of view. One field of view images simultaneously at 3.6 and 5.8 μm and the other at 4.5 and 8.0 μm via dichroic beamsplitters. All four detector arrays are 256 × 256 with 1.2″ square pixels. IRAC is described in detail in Chapter 6.

Table 2.1: Summary of Spitzer Characteristics

Aperture (diameter)	85 cm
Orbit	Solar (Earth-trailing)
Cryogenic Lifetime	5.5 years (approx)
Wavelength Coverage (passband centers)	3.6 - 160 μm (imaging) 5.3 - 40 μm (spectroscopy) 55 - 95 μm (spectral energy distribution)
Diffraction Limit	5.5 μm
Image Size	1.5″ at 6.5 μm
Pointing Stability (1σ, 200s, when using star tracker)	<0.1″
As commanded pointing accuracy (1σ radial)	<0.5″
Pointing reconstruction (required)	<1.0″
Field of View (of imaging arrays)	~ 5′×5′ (each band)[1]
Telescope Minimum Temperature	5.6 K
Maximum Tracking Rate	1.0″/ sec
Time to slew over ~90º	~8 minutes

The InfraRed Spectrograph (IRS) – James R. Houck, Cornell University, PI
IRS performs both low and high-resolution spectroscopy. Low-resolution, long slit spectra ($\lambda / \Delta\lambda$ =64–128) can be obtained from 5.2 to 38.0 μm. High-resolution spectra ($\lambda / \Delta\lambda$ ~600) in Echelle mode can be obtained from 9.9 to 37.2 μm. The spectrograph consists of four modules, each of which is built around a 128×128 pixel array. One of the modules incorporates a peak-up function that can be used in locating and positioning sources on any of the four spectrometer slits with sub-arcsecond precision. The IRS Peak-Up array has 1.8″ square pixels and a field of view of 1′ by 1.2′. Two filters are available for use with an IRS peak-up (and imaging), covering 13.5–18.5 μm and 18.5–26 μm, respectively. IRS is described in detail in Chapter 7.

The Multiband Imaging Photometer for Spitzer (MIPS) – George H. Rieke, University of Arizona, PI
MIPS is designed to provide photometry and super resolution imaging, as well as efficient mapping capabilities, in three wavelength bands centered near 24, 70 and 160 μm. The array materials, sizes and pixel scales vary; they are given in Table 2.2. MIPS is also capable of low-resolution spectroscopy ($\lambda / \Delta\lambda$ ~15–25) over the wavelength range 55–95 μm and a Total Power Mode for measuring absolute sky brightness. MIPS is described in detail in Chapter 8.

---

[1] At 160 μm, a single field of view is 0.53′×5.33′. See MIPS chapter for more information.

2.4 Spitzer Science Operations Overview

The Spitzer Science Center (SSC) conducts the science operations for Spitzer, and is charged with 1) acting as an interface and advocate for users, 2) capturing and conducting the science program efficiently, 3) producing and securing the Spitzer science legacy, and 4) conducting public and scientific outreach for the Spitzer program. The SSC is located on the campus of the California Institute of Technology in Pasadena, California, USA.

In carrying out its charter, the SSC issues annual Calls for Proposals, organizes science and technical reviews, selects approved programs based on a Time Allocation Committee review, and administers data analysis awards. The SSC provides tools for detailed planning of Spitzer observations and proposal submission, and offers Science User Support services. In addition, the SSC schedules observations on the telescope, provides basic (pipeline) science data processing and data quality assessment, and creates a data archive that has been publicly accessible since May 2004.

2.4.1 Science Instrument Operations

Spitzer's payload is capable, but simple. Among the three science instruments there are only two moving parts: the shutter on IRAC (current operational plans do not include the use of the shutter) and the scan mirror on MIPS. The observatory schedule is organized into instrument campaigns. Only one of the three instruments can be powered on at any one time, but all of the apertures of that instrument are available whenever the instrument is on (with the exception that only one of the four IRS modules can be read out at a time). As shown in Figure 2.1, the apertures point toward different portions of the sky at the same time; switching apertures and staying on the same target in general involves re-pointing the telescope. A few things to notice are that the IRS slits cover a wide range of orientations and that the MIPS fields of view are larger than the array size; this is because they show the range covered by the motion of the MIPS scan mirror.   (See also Figure 4.5 and discussion in Section 4.3.)

A central concept in Spitzer science operations is that of the Astronomical Observation Templates (AOTs). The three science instruments are operated in eight discrete observing modes that offer the observer a selected number of choices in configuring and operating the instrument. These observing modes are known as AOTs. The AOT concept and the eight Spitzer AOTs are described in Chapter 5, and in greater detail in the instrument-specific Chapters.

The SSC plans and executes calibration activities to maintain the instrument calibration. In general, observers should not need to supplement the instrument calibration, but may do so by incorporating observations for that purpose into their (proposed) observing program. The instrument calibration for each instrument is discussed in its relevant chapter. Calibration across instruments is no longer described in this document. Information about cross-calibration can be obtained from our website at http://ssc.spitzer.caltech.edu/obs/xcalib.html.

Table 2.2: Spitzer Instrumentation Summary (NB: limits include confusion)

	λ (μm)	Array Type	λ/Δλ	Field of View	Pixel Size (arcsec)	Sensitivity [2] (μJy) (5 σ in 500 sec, incl. confusion)
IRAC	3.6	InSb	4.7	5.21′×5.21′	1.2	1.6 (3.4)[3]
	4.5	InSb	4.4	5.18′×5.18′	1.2	3.1 (4.3)
	5.8	Si:As(IBC)	4.0	5.21′×5.21′	1.2	20.8 (21)
	8.0	Si:As(IBC)	2.8	5.21′×5.21′	1.2	26.9 (27)
IRS	5.2–14.7	Si:As(IBC)	64–128	3.7″×57″	1.8	250 [4]
	13.5–18.5 18.5–26	Si:As(IBC) Peak-Up [5]	~3	54″×80″	1.8	116 80
	9.9–19.5	Si:As(IBC)	~600	4.7″×11.3″	2.3	1.2×10^{-18} W/m^2
	14.3–35.1	Si:Sb(IBC)	64–128	10.6″×168″	5.1	1500
	18.9–37.0	Si:Sb(IBC)	~600	11.1″×22.3″	4.5	2×10^{-18} W/m^2
MIPS	24	Si:As(IBC)	5	5.4′×5.4′	2.55	110 [6]
	70	Ge:Ga	4	2.7′×1.4′ 5.2′×2.6′	5.20 9.98	14.4 mJy[7] 7.2 mJy
	55-95 [8]	Ge:Ga	15–25	0.32′×3.8′	10.1	57, 100, 307 mJy (@60, 70, 90 μm)
	160	Ge:Ga (Stressed)	5	0.53′×5.33′	16×18	29 (40) mJy[9]

---

[2] Sensitivities given here are for point sources, and are only representative. See Chapters 6-8 for more detail.

[3] IRAC sensitivity is given for intermediate background. The first number in each case is without confusion, and the second number (in parentheses) includes confusion.

[4] IRS sensitivity is given for low background at high ecliptic latitude. Note that for IRS, sensitivity is a strong function of wavelength.

[5] For recommended flux density range for peak-up target, please refer to IRS chapter (chapter 7).

[6] MIPS sensitivity is given for low background.

[7] 70 μm can be confusion limited; see the MIPS chapter for details.

[8] Because of a bad readout at one end of the slit, the spectral coverage for 4 columns of the array is reduced to about 65-95 microns.

[9] 160 μm is often confusion limited; 29 mJy refers to no confusion, and 40 mJy refers to the estimated confusion limit.

Nominal Field-of-View Locations
Projected onto the Sky

Note: This is not an engineering diagram.
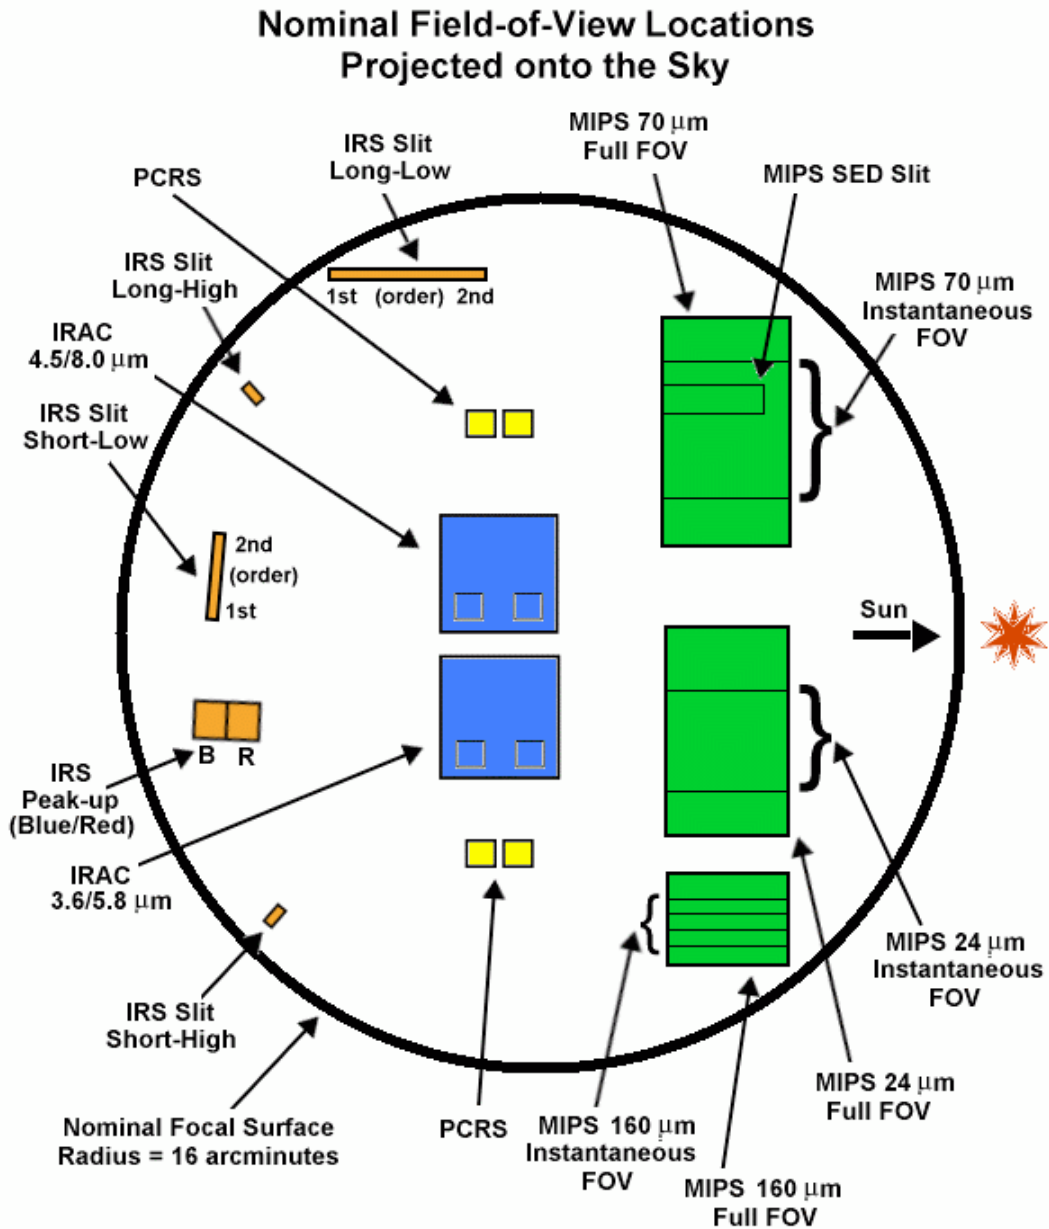
Figure 2.1: Science instrument apertures projected onto the sky. Because of the optical inversion in this projection, the section of sky closest to the projected Sun is on the MIPS side of the focal plane, e.g. to the right in this view. Because the spacecraft does not rotate about the line of sight, this vector is fixed relative to the focal plane on the sky. The IRAC sub-array fields are shown by the small boxes in the lower corners of both IRAC arrays. (The 8.0 and 5.8 µm sub-arrays are on the right and the 4.5 and 3.6 µm sub-arrays are on the left.) Note that for figure clarity, the widths of the IRS slits as shown are rendered substantially larger than their actual scale. See Chapters 6–8 for more information on the instruments.

2.4.2 The Life Cycle of a Spitzer Observation
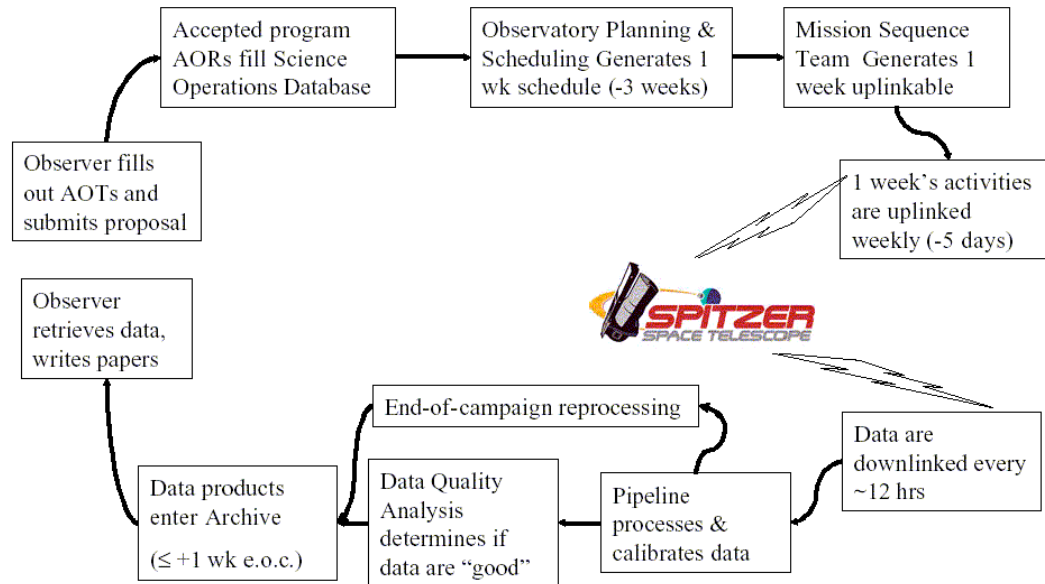
Figure 2.2: The life cycle of a Spitzer observation. The detailed definition of observations will occur at the time of each annual proposal solicitation.

Figure 2.2 shows the basic stages through which a Spitzer observation passes. These processes are described in more detail in Chapter 5. The process of defining a Spitzer observation starts with the choice of an astronomical target and the selection of one of the eight observing modes (AOTs). An SSC-supplied software tool called Spot (see section 5.3.1) is used to enter the target information and the observation details. After the AOT is chosen, Spot is used to enter all the parameters (e.g., integration time, choice of modules, etc.) needed to fully specify the observation. This process is called "filling out the AOT front-end." Spot also provides wall-clock time estimates for the total duration of observations. Spot does not allow observers to create impossible observations (however, observer-supplied constraints can render an observation unschedulable; see section 5.5.3)

An observation which has been fully defined by supplying parameter values for an AOT is known as an Astronomical Observation Request (AOR), and is the basic scheduling unit for Spitzer. After the Time Allocation process has been completed, the AORs submitted for all of the approved Spitzer observations are loaded into a database known as the Science Operations Database (SODB), and are stored there in the form of a set of specific parameters and values. These parameters are used by the software AOR/IER Resource Estimator (AIRE) to expand the set of parameters into a set of instrument and spacecraft commands which execute the observation on board Spitzer, as well as to provide extremely high fidelity estimates of execution time for the scheduling process. Observers also access AIRE through Spot for their time estimates. The Spitzer planning and

scheduling process produces observing schedules based on these resource estimates. After the content of the schedules has been approved and finalized, AIRE takes the AOR parameters and produces command product files which are then processed by the JPL Mission Sequence Team and finally uplinked and executed by Spitzer. *It is important to realize that the process of creating commands to carry out an observation based on the AOR parameters is done by software, not by support astronomers at the SSC.*

After an observation has been scheduled and carried out, the resulting data are pipeline-processed, undergo a brief quality-checking process, and are placed in the Spitzer science archive and made available to the observer. For Legacy Science projects, the processed data are also placed in the Spitzer public archive which is available for the use by the whole community. As of the opening of the public archive in May 2004, placement of these data into the public archive is made at the same time as they are passed on to the Legacy Science teams.

2.5 SSC Science User Support Services

The Spitzer Science Center acts as the interface between the astronomical community and the observatory. The SSC offers support at all stages of Spitzer use, including proposal preparation, detailed observation design, and data analysis. The SSC provides this support primarily through supplying documentation, providing science user tools, and operating an online Helpdesk.

2.5.1 Documentation

The primary form of science user support is through the dissemination of information, including this Observer's Manual (SOM). In addition to the SOM, the SSC offers a comprehensive set of web pages, which includes further technical documentation, software downloads, Spitzer Frequently Asked Questions (FAQ) pages, observation scheduling status, frequently updated news, instrument status, performance reports plus tools, and information to assist with proposal preparation.

2.5.2 Helpdesk

The SSC Helpdesk service (help@spitzer.caltech.edu) is the primary interface to the Spitzer Science Center for the community. The Helpdesk provides answers to technical or policy questions, and is available to address concerns about any facet of Spitzer observing. The Helpdesk also provides support for special needs such as the activation of Target of Opportunity (ToO) proposals, requesting modifications to Spitzer observations, and the submission of proposals for Director's Discretionary Time (DDT).

Once a Helpdesk question has been received, it is directed to an appropriate expert. An observer can expect a reply within two business days (except for ToO activation requests, for which an initial response within about 2 hours can be expected). Particularly complex requests may take up to 5 days for a response. On the other hand, easy questions can result in a nearly instantaneous response.

(There is no auto-reply on our system.) Note that as the proposal deadline approaches, the turnaround times for an SSC response (particularly for definitive answers to complex technical questions) will likely increase. It is the responsibility of proposers to take this reality into consideration when submitting queries shortly before the proposal submission deadline.

2.5.3 Proposal Kit

The online Proposal Kit provides all of the information necessary for the prospective General Observer (GO), Archival Research (AR), or Theoretical Research (TR) investigator. It can be found on the SSC website at http://ssc.spitzer.caltech.edu/propkit/.

The Proposal Kit also includes instructions for downloading and installing the Spot software on the user's machine. Spot, discussed more in section 5.3.1 below, is used to plan and prepare Spitzer observations, and to submit GO, AR, and TR proposals electronically to the SSC. *Prospective GO/AR/TR investigators are recommended to download Spot and to start planning their Spitzer observing programs well before the proposal submission deadline.*

In addition, the Proposal Kit includes many kinds of documentation, including this manual, the current Call for Proposals, Spitzer Observing Rules, an Observation Planning Cookbook, Proposal Submission Guidelines, and the proposal template files, guidelines pertinent to the calculation of the special overhead burdens applied to Targets of Opportunity, targets requiring rapid instrument changes, and targets with late ephemeris changes, etc.

A separate section of the online Proposal Kit contains information specifically designed for Spitzer Solar System researchers. It includes asteroid count estimates, ephemeris files currently integrated into Spot, NAIF name identifications, and tips for using the Horizon database supported by the Solar System Dynamics Group at the Jet Propulsion Laboratory.

The SSC website also contains an "Infrared Compendium," a resource for professional astronomers new to infrared astronomy. It is a catalog of selected archetypal objects and their expected fluxes in the infrared, along with some overview material that discusses issues relevant to infrared observers, such as resolution and background sources, along with specific applications to Spitzer.

Finally, the Proposal Kit contains a multitude of links to useful general-purpose astronomical tools (such as coordinate conversion), observation planning tools (IRSky, Skyview), astronomical databases (IRSA, NED, SIMBAD), and archived datasets (2MASS, ISO, DSS).

2.5.4 Data Products

Once data are collected, the SSC creates and delivers pipeline-processed, calibrated data products to the observer and to the Spitzer science archive. Three levels of processed data are created: Raw data, Basic Calibrated Data (BCD) and

Ensemble data (also known as "post-BCD data"). Raw data are unprocessed data, which have been packaged into FITS format. BCDs are single-frame FITS format data products, which have been processed to remove instrumental signatures, and which are calibrated in scientifically useful units. Ensemble data are higher-level products, and may include mosaicking, co-addition, or spectrum extraction. The data products are described in more detail in the instrument-specific chapters and in the corresponding Data Handbooks, available on the SSC website.

In general, the final calibration for an instrument campaign, which is typically about one–two weeks in duration, is based upon calibration observations taken at the beginning and end of the campaign, so pipeline data products are typically available a few days to weeks after the end of an instrument campaign.

The SSC validates the pipeline processing for each observing mode prior to the release of data. A basic quality assessment is performed on all Spitzer data before they are delivered to the observer. Quality assessment information is also available to the observer through the archive interface.

In order to obtain data, investigators should use the Leopard interface, which is not only very similar to Spot (see above), but also downloaded and installed as part of the Spitzer Pride, the same package as Spot. The Leopard User's Guide is also available on the SSC website. Certain "Popular Products" (e.g. FLS, original Legacy Program data, etc.) are available through the Leopard interface and through the SSC website.

After obtaining data, investigators may be interested in the Data Handbooks. Data Handbooks are available for all three instruments, linked off their respective pages at the SSC.

3 Observatory Description and Constraints

This section addresses the Spitzer Space Telescope design and the observational constraints that arise from that design. The Spitzer Space Telescope consists of five basic elements: the cryogenic telescope assembly (CTA), the spacecraft, and the three science instruments. In this section the spacecraft and CTA design (not including the telescope design, see Chapter 4) are discussed. The design and performance of the pointing and control system (PCS) are also discussed.

A brief, high-level summary of the mission for astronomers appears in the ApJS Spitzer Special Issue, specifically the paper by Werner et al. (2004, ApJS, 154, 1), entitled "The Spitzer Space Telescope Mission." A copy of this paper is available on the SSC website.

3.1 Observatory Design and Operations Concept

Spitzer is unique in many ways; the cryostat and mission design allow a very long cryogenic lifetime (~5.5 years) with a relatively small amount of cryogen (~360 liters of superfluid helium)[10]. The helium bath temperature is maintained at 1.24 K, the telescope temperature can be ≤ 6 K, and the outer shell temperature is approximately 34 K. Some of Spitzer's unique features are its Earth-trailing solar orbit and the fact that the telescope was launched at ambient temperature. This "warm launch" architecture allows Spitzer to achieve its lifetime goal with a smaller, lighter cryostat than would otherwise be required.

A basic external view of Spitzer is shown in Figure 3.1. The spacecraft provides structural support, pointing control and telecommunications and command/data handling for the entire observatory. The spacecraft, including the spacecraft bus, the solar panel and PCS (including the Pointing Control Reference Sensors, or PCRSs, which are located in the focal plane) were provided by Lockheed Martin Space Systems Company.

The CTA, shown in Figure 3.2, consists of the telescope, the superfluid helium cryostat, the outer shell group and the Multi-Instrument Chamber (MIC), which hosts the cold portions of the IRAC, IRS, and MIPS science instruments and the PCRS. The telescope assembly, including the primary and secondary mirrors, metering tower, mounting bulkhead (all made of beryllium) and the focus mechanism, is mounted and thermally connected to the cryostat vacuum shell. The barrel baffle of the telescope assembly is separately attached to the vacuum shell at its flange. The MIC, an aluminum enclosure containing the instrument cold assemblies, is mounted on top of the helium tank. The vapor-cooled shields, the vacuum shell and the outer shields are connected to the cryostat by a series of low connectivity struts. The CTA was provided by Ball Aerospace and Technologies Corporation.

The spacecraft and CTA (and science instruments) together comprise the observatory, which is ~4 m tall, ~2 m in diameter, and has a mass of ~900 kg.

---

[10] For comparison, IRAS had a 10-month lifetime using 560 liters of liquid helium, and ISO had a 28-month lifetime using 2140 liters of cryogen.
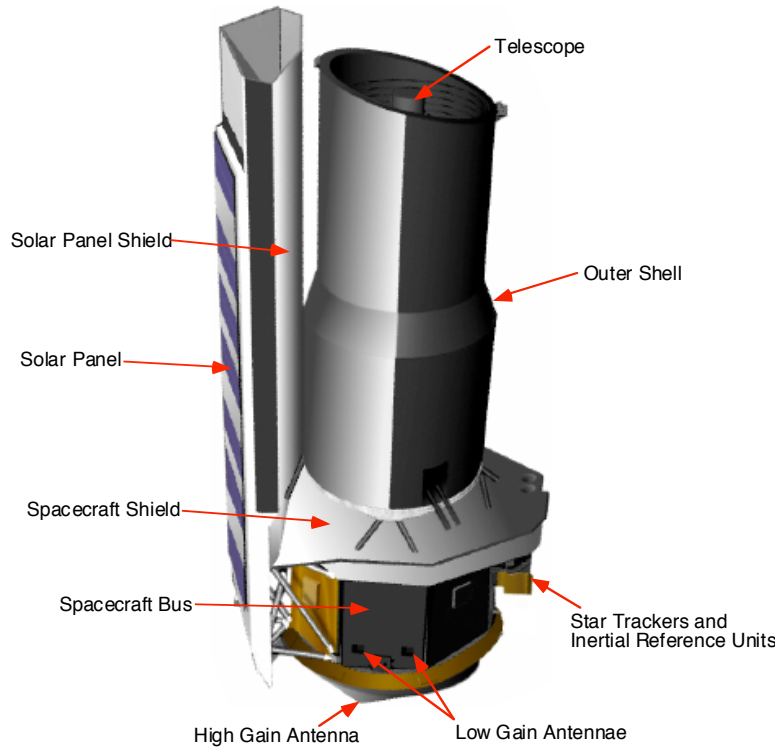
Figure 3.1: Basic external view of Spitzer.

The observatory coordinate system XYZ (shown in Figure 3.3) is an orthogonal right-hand body-fixed frame of reference. The X axis passes through the geometric center of the top surface of the spacecraft, is parallel to the CTA optical axis (which passes through the primary and secondary mirror vertices), and is positive looking out of the telescope. The Z-axis intersects the line forming the apex of the two surfaces of the solar panel. The Y-axis completes the right hand orthogonal frame. The X-axis origin is defined such that the on-axis point between the CTA support truss and the spacecraft bus mounting surface is located at X = +200 cm, in order to maintain positive X values throughout the observatory. The Sun always lies within 2º of the XY plane (i.e., the roll angle is constrained to ±2º).

Spitzer is operated autonomously for moderately long periods of time (12 to 24 hours) interspersed with short periods (~20–80 minutes) of ground contact. (One contact per day is always a minimum of 40–50 minutes long.) In extreme circumstances, Spitzer is designed to survive for up to a week with no ground contact at all.

During routine science operations, Spitzer typically executes a pre-planned week-long schedule of science observations, calibrations and routine engineering activities, which has been uploaded in advance and stored on board. This "master sequence" might typically have 14 different 12–24 hour Periods of Autonomous Operation (PAOs), containing observations and calibration activities, each

followed by a 30-60 minute period spent re-orienting the spacecraft for downlink and transmitting the data to the ground. For IRS campaigns, we routinely have a 24-hour PAO with one downlink lasting longer than 40 minutes. After the downlink, Spitzer returns to the pre-planned sequence of observations and calibrations. Because efficient communication with the ground requires use of the high-gain antenna mounted on the bottom of the spacecraft (Figure 3.1), executing a downlink of the collected data requires slewing the spacecraft to orient the high gain antenna toward one of the Deep Space Network (DSN) stations on Earth. Any of the three DSN sites (Canberra, Madrid and Goldstone) can be used when visible. During the time that data are being transmitted to the ground, no science data can be collected.
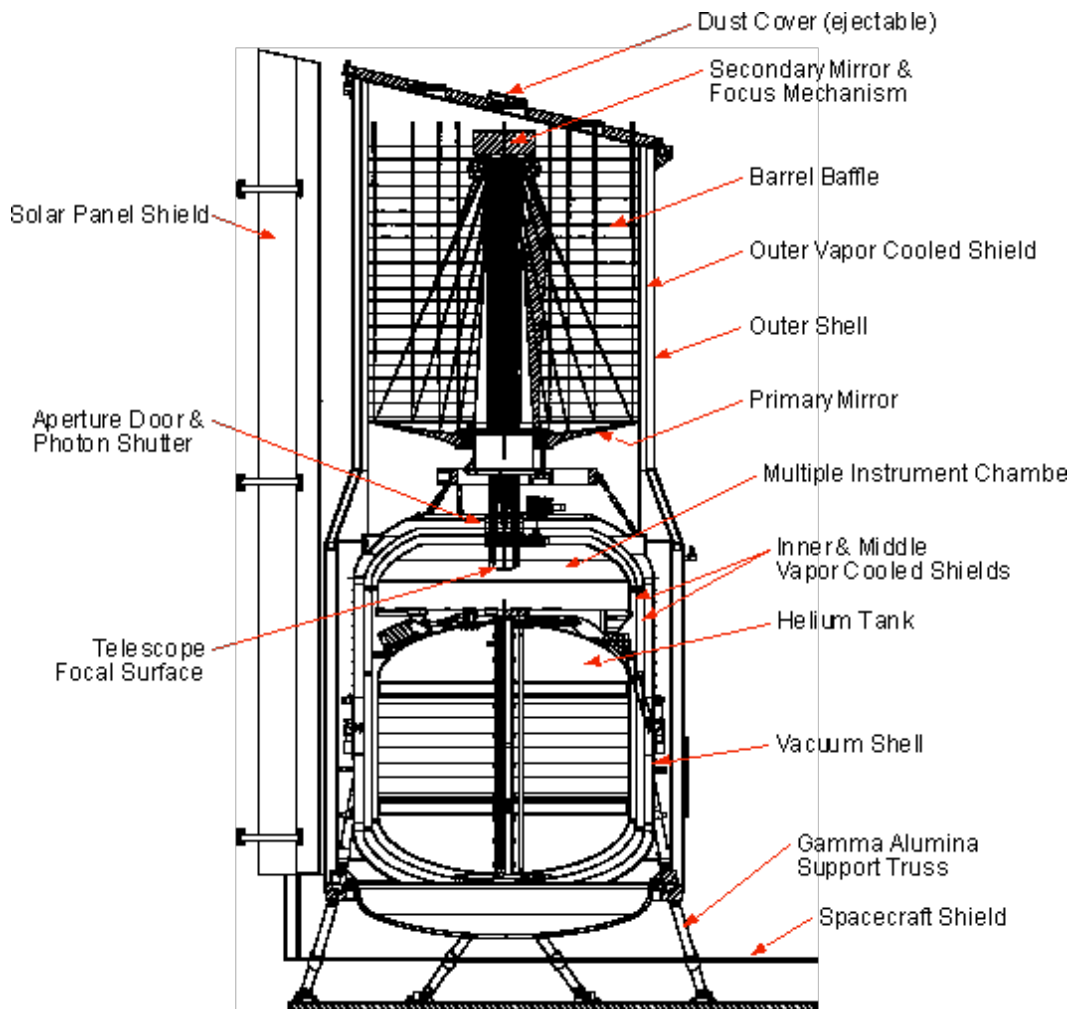
Figure 3.2: Cryogenic Telescope Assembly.

S/C-to-CTA interface plane at station X = 200 cm
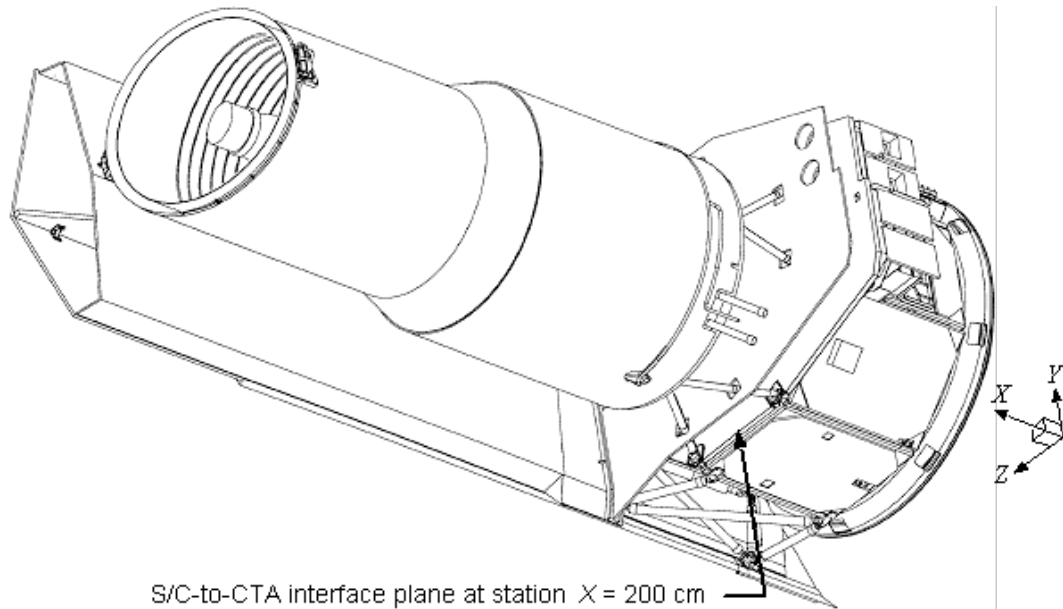
Figure 3.3: Observatory Coordinate System.

If a downlink opportunity is missed (e.g., due to a failure at the ground station), Spitzer preserves the collected science data, continues collecting new science data as scheduled, and attempts to downlink it at the next scheduled opportunity. Once a missed downlink opportunity has occurred, the spacecraft sequence is modified to provide additional downlink time within the next few opportunities if necessary. No data are lost or overwritten, unless the risk of overfilling the entire mass memory (16 Gbits) exists. (See also section 3.1.2 below.)

3.1.1 The Warm Launch and Telescope Temperature Management

One of the unique aspects of Spitzer is the warm launch concept. The telescope was at ambient temperature at the time of launch and gradually reached its lowest operating temperature (≤ 6 K) approximately 45 days after launch, through passive cooling and by vapor vented from the cryostat. The science instruments are in contact with the 1.4 K helium bath at all times and were thus cold at the time of launch. The on-orbit performance confirms the soundness of the Spitzer thermal design principles.

Because the telescope is cooled by vapor vented from the cryostat, the telescope temperature depends upon how much power is dissipated in the cryostat. Because the cryogenic design and performance result in near-zero parasitic heat loads into the cryostat, the power entering the helium bath is essentially that of the instrument in use. Due to variations in the power dissipation of the instruments, the telescope temperature would naturally vary from $\leq$ 6 K to ~12 K. A "make-up" heater is available to add additional heat when needed to lower the telescope temperature. Since very low background is required for 160 μm observations, make-up heater operations are always planned to cool the telescope to ≤ 6 K and

maintain it at that temperature during MIPS campaigns that are obtaining 160 μm observations. This is accomplished using a heat pulse that has been optimized to use the least amount of cryogen, while still guaranteeing a cold telescope for MIPS; it uses considerably less cryogen than holding the temperature constant all the time. The heat pulse is typically ~10 hours in duration and takes place about 10 days prior to the start of the MIPS campaign. The details of pulse size and duration are optimized to minimize cryogen usage and maximize scheduling efficiency.

In order to maximize Spitzer's cryogenic lifetime, MIPS campaigns will be organized into "warm" (telescope cooled to ~8.5 K) or "cold" (~5.5 K) campaigns. The cold campaigns, in which 160 μm observations can be executed, will likely be scheduled as these observations dictate, yet still enabling observations of the full sky. If MIPS scan AORs are requested and the 160 μm observations are required for the science, the "160 μm required" checkbox must be selected in the AORs. AORs requiring 160 μm will be segregated and scheduled only in cold (~5.5 K) MIPS campaigns. MIPS observations not requiring 160 μm will be scheduled in warm (~8.5 K) or cold (~5.5 K) campaigns. Programs requiring 160 μm data will be more difficult to schedule, given that not all MIPS campaigns in Cycle-5 will be "cold."

Cycle-5 MIPS scan observers will need to specifically justify their need for the 160 μm array. Proposers should carefully consider the science objectives of their program before requesting 160 μm data. See the Call for Proposals for more details.

3.1.1.1 Impact on Observers

In general the variation of telescope temperature need not be taken into account by observers during observation planning. For Target of Opportunity (ToO) observations that require optimal 160 μm data from MIPS, an additional scheduling delay may be incurred by the need to reduce the telescope temperature prior to observation. It is possible for the observer to specify if highest-quality 160 μm data are required when observing with MIPS; this is used to ensure the observation is scheduled when the telescope is at ≤ 6 K. Cycle-5 MIPS scan observers will need to specifically justify their need for the 160 μm array; see the Call for Proposals for more details.

3.1.2 Data Storage

The Spitzer spacecraft includes the command and data handling subsystem (C&DH) which shares the flight computer (a RAD 6000) with the pointing control subsystem (PCS). The C&DH validates and executes either previously stored or real time commands, receives and compresses data, and writes the compressed data into the mass memory. The C&DH also provides a stable clock to correlate data and events.

During normal Spitzer science operations, the observatory collects, compresses, and stores 12 to 24 hours' worth of science data prior to downlinking it. Spitzer has 8 Gbits of solid state memory available for each of two redundant flight computers, for a total of 16 Gbits of storage (for both science and engineering data). During each 12 to 24-hour PAO, up to 6 Gbits of the memory are filled. Enough additional storage capacity is left unused during that 12 to 24-hour PAO to permit missing a downlink opportunity (e.g., due to a problem at the ground station) and continuing to observe without overwriting previously collected science data. Since the full mass memory (16 Gbits) is available to both redundant computers, in general multiple passes would have to be missed before the risk of losing science data that had not been transmitted to the ground becomes significant. However, e.g., MIPS sometimes generates enough data that a single missed pass could become important.

3.1.3 Data Transmission

Spitzer has two antenna systems for data uplink and downlink. The high gain antenna (HGA) supported a maximum downlink rate of 2.2 Mbit/s for the first 2.6 years of the Spitzer mission. Since Spitzer drifts away from the Earth at about 0.12 AU per year, the data transmission rate on the HGA may have to decrease as a function of time, and late in the mission, we may have to change our downlink strategy to accommodate this. The two low gain antennae (LGA) give wide angle coverage, but only support downlink at 44 kbit/s at the beginning of routine observing, decreasing to 40 kbit/s after 3 years. The LGA are used in safe mode, but are not normally used to transmit science data during routine science operations.

Because of the location of the HGA, it is necessary to stop observing and point the telescope in the anti-Earth direction to downlink the science data. Depending upon the details of the observations that are scheduled, Spitzer produces ~1 to 6 Gbits of (compressed) science data during 12–24 hours of observing. Initially, Spitzer collects data for about 12–24 hours and then spends 0.5–1 hour downlinking the data. The scheduling system predicts the compressed data volume that is generated during each downlink and schedules the downlink contact time accordingly. The downlink periods are also used for some spacecraft maintenance activities (such as dumping angular momentum) and for uplinking commands. If low volumes of data are being generated, longer periods of time between downlinks may be used.

At each downlink opportunity, Spitzer attempts to transmit all the data in the mass memory. Any data that the ground has not confirmed as received will be retransmitted at the next downlink pass. It is anticipated that some data could be missed at any pass, and it may take several passes before the ground receives all the science data for a given observation.

Whenever the HGA is used for downlink, it is also possible to uplink commands and files to Spitzer. Normally, up to ~5 communication periods per week are used to send up all the sequences and information needed to support the next

upcoming week of observations. No communications with Spitzer are planned outside the scheduled downlink sessions.

3.1.3.1 Impact on Observers

The downlink strategy has little impact on observation planning with Spitzer. It is one component in determining the maximum length of an AOR, but the necessity for instrument maintenance operations and pointing system calibrations on shorter than 12 hour time scales constrains the maximum AOR length more tightly than the downlink schedule[11]. A 12 to 24 hour interval between downlinks does imply that it can be a day or longer before science data are available on the ground after an observation is complete.

The uplink strategy also does not directly affect observation planning, but it drives the time scale on which changes to the stored science schedule and its contents can be made (e.g., for a ToO).

3.2 Sky Visibility

3.2.1 Solar Orbit

An important innovation enabling Spitzer to accomplish ambitious scientific goals at a modest cost is its Earth-trailing heliocentric orbit (shown in

Figure **3.4**), in which the observatory drifts away from Earth at the rate of ~ 0.12 AU/year. This will substantially help prolong the coolant lifetime by supporting an operating regime in which most of the cryogen is used to take up the power dissipated by the detector arrays, rather than lost to parasitic heat loads. In addition, this orbit has less-constrained visibility, compared to what would have been the case in a near-Earth orbit, allowing all parts of the sky to be visible for at least two extended periods each year and some zones to be visible continuously.

3.2.2 Pointing Constraints

Spitzer's view of the sky is limited by two hard pointing constraints, illustrated in Figure 3.5 and (in a different fashion) in Figure 3.6:

- The angle between the boresight and the direction of the Sun may never be less than 82.5º. (NB: this was updated in early 2004.)
- The angle between the boresight and the direction of the Sun may never exceed 120º.

The area defined by these hard constraints is called the Operational Pointing Zone (OPZ). In addition, some bright objects (such as the Earth) are normally avoided, because they would degrade the quality of the observation, due to direct exposure or stray light, but this is not a strict constraint (see section 3.2.6 and Appendix

---

[11] Currently, the maximum length of an AOR is 8 hours for IRAC and IRS, and 3 hours for MIPS. Please note that this is the maximum length for an isolated AOR. In general, there would not be several contiguous periods of this length available.

10.2). Note that the definition of the OPZ precludes observing Mercury or Venus with Spitzer. A second-order effect on the OPZ is provided by the limited roll angle of Spitzer around the Y-axis (see Figure 3.3), which is just ±2º.

Figure 3.6 shows the actual pointings of the observatory for one day in the life of Spitzer. Note the locations of the OPZ and how it constrains where we actually pointed on that day.
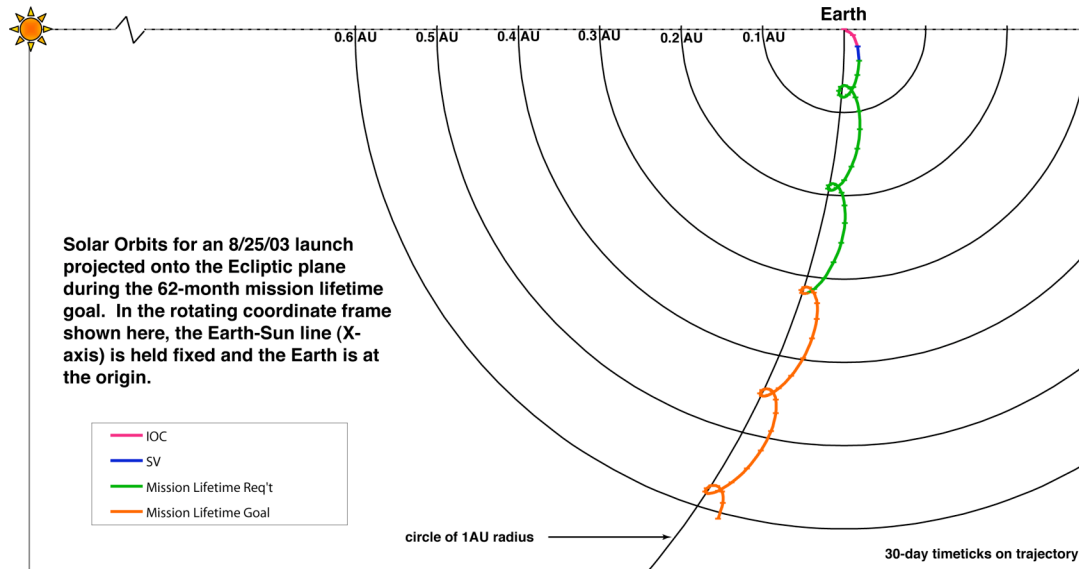
Figure 3.4: Spitzer's solar orbit projected onto the ecliptic plane and viewed from ecliptic North. In the rotating frame, the Earth is at the origin and the Earth-Sun line is defined as the X-axis. "Loops" and "kinks" in the trajectory occur at approximately 1-year intervals when Spitzer is at perihelion. Spitzer's orbit is also slightly inclined with respect to the ecliptic.

3.2.3 Viewing Periods

The amount of time during the year any particular target is visible depends primarily on the absolute value of its ecliptic latitude (Figure 3.7). As seen in Figure 3.5, Spitzer's instantaneous window of visibility on the celestial sky forms an annulus (the OPZ), perpendicular to the ecliptic plane, of ~40º width and symmetrical with respect to the Sun. This annulus rotates with the Sun over the period of a year; the edges of the OPZ move along the ecliptic at about 1º/day as Spitzer orbits the Sun. For an object near the ecliptic plane, the length of the visibility period is ~40 days twice a year (modulo periods when undesired bright moving objects are also present and bright object avoidance has been selected; see section 3.2.6). The visibility periods increase to ~60 days twice/yr at an ecliptic latitude of ~45º, ~100 days twice/yr at latitudes ~ 60º, then becoming a single long window ~250 days long for latitudes near 60-70º, finally reaching constant viewing near the ecliptic poles. About a third of the sky is visible to Spitzer at any time. Figure 3.8 illustrates how the total number of days of visibility varies over the sky in three coordinate systems.
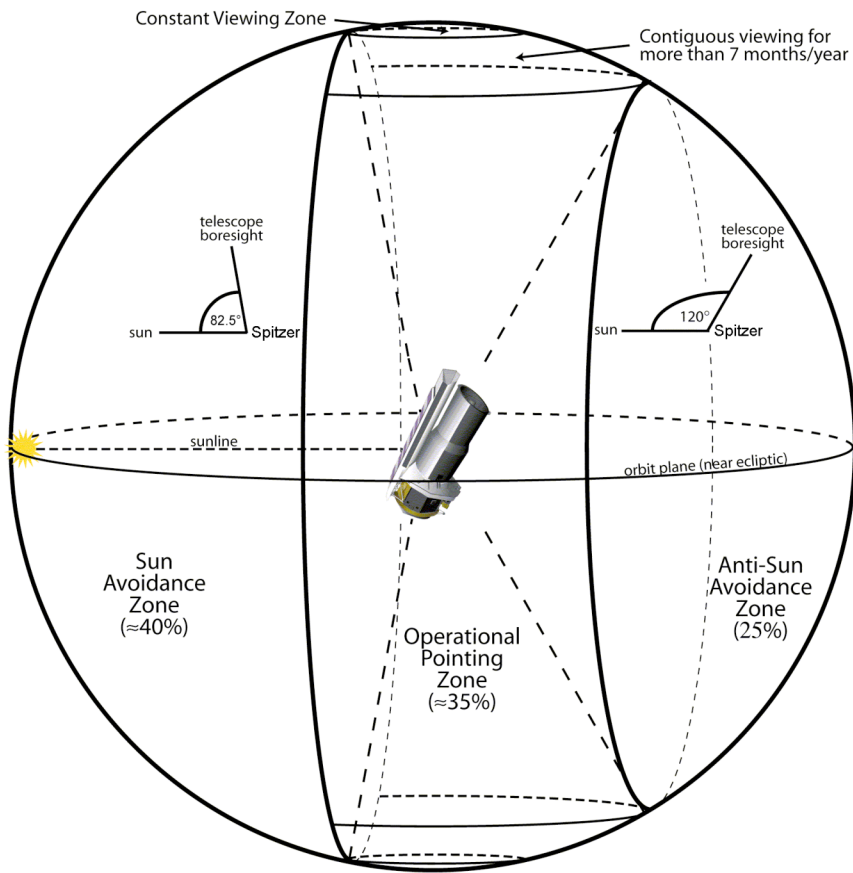
Figure 3.5: The main geometric observing constraints form an area called the Operational Pointing Zone (OPZ).
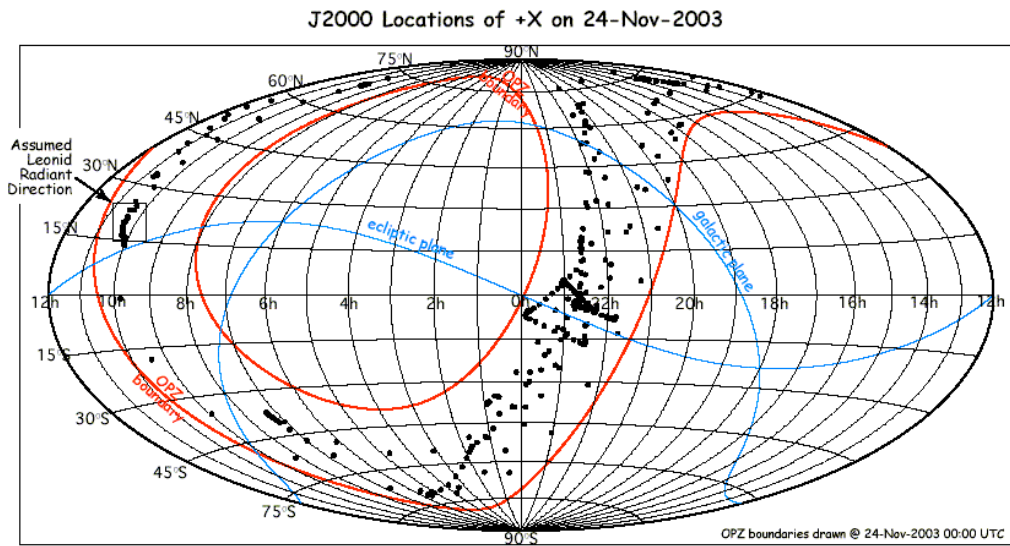
Figure 3.6: OPZ boundaries for 24 Nov 03 (0h UTC), with dots representing the actual locations of the telescope boresight for the subsequent 24 hours.
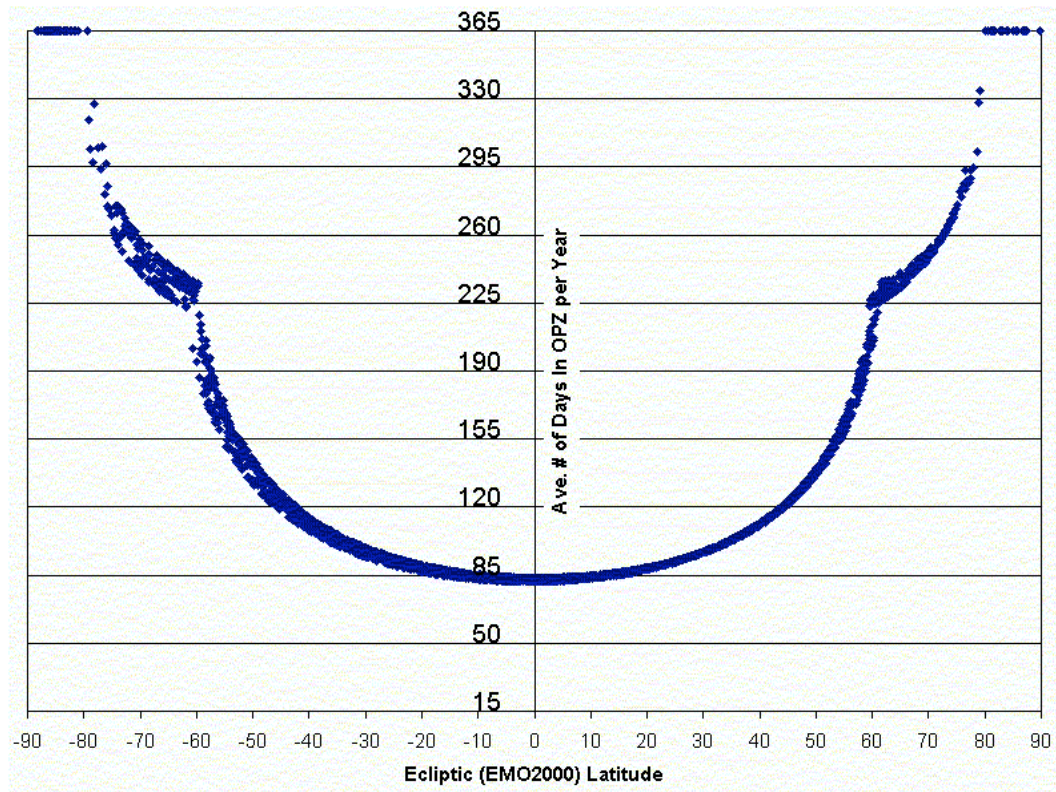
Figure 3.7: Variation of length of visibility period as a function of ecliptic latitude for all of the targets in the April 2003 ROC. (This figure is provided as indicative of the general concepts, despite the fact that the ROC has changed substantially since April 2003.)
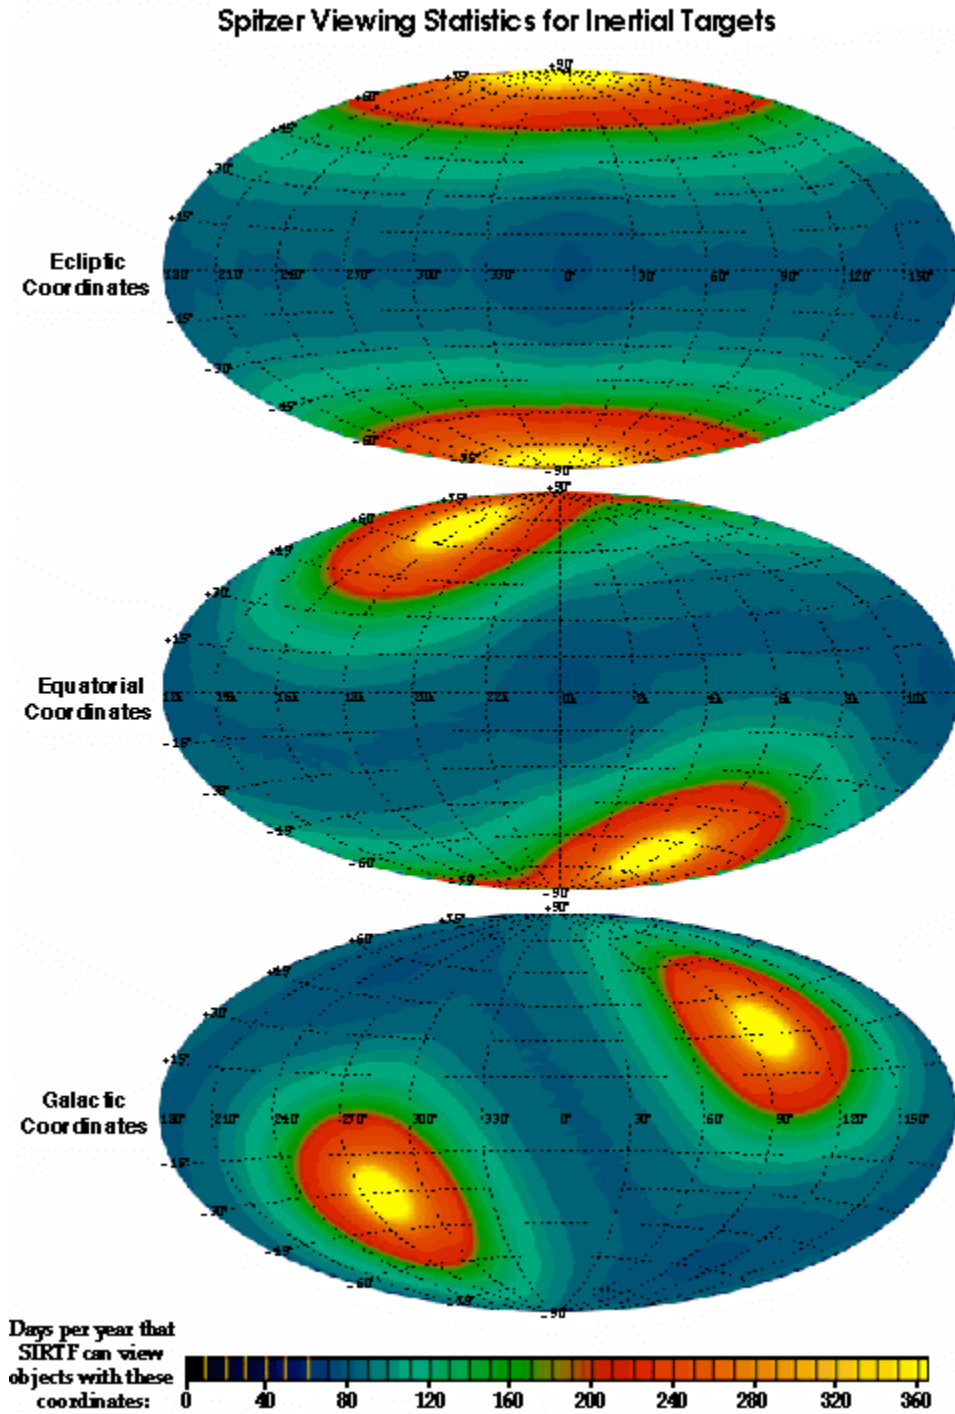
Figure 3.8: Total days of visibility per year in equatorial, ecliptic and galactic projections.

3.2.4 How Visibility Evolves with Time

Figure 3.9 illustrates the zone of visibility in equatorial coordinates for four specific dates during the first year of operations. Note how the sky passes into and out of visibility as the calendar date advances. Each set of contours shows the available and forbidden zones on a particular date (in this plot, 1 Sep 2003, 1 Dec 2003, 1 Mar 2004, and 1 Jun 2004). On each date, there is a forbidden zone in the anti-Sun direction, as well as a second, more extended-looking forbidden region on the side of the spacecraft towards the Sun; the Sun-ward forbidden region is in the center of the plot on 1 Sep and on the edges for the 1 Mar plot. Note that even targets at high declinations may fall in the forbidden zones part of the time, although objects at extreme declinations are generally visible for most of the year.

Although these plots are accurate, observers should always use Spot to determine visibility of a target on a particular day (see Chapter 5).
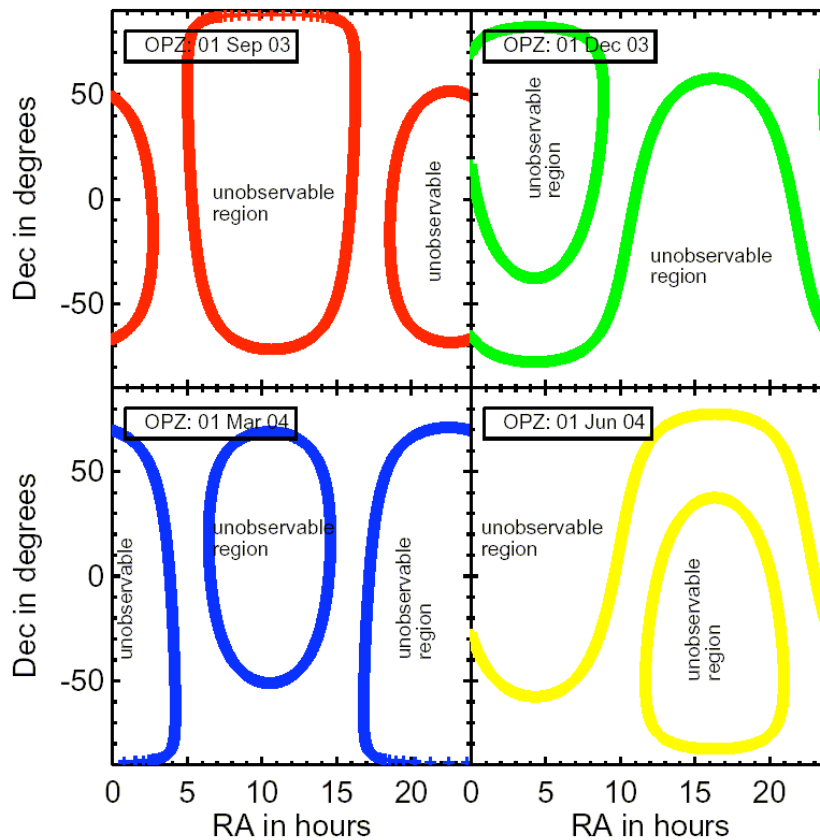
Figure 3.9: Example of time evolution of visibility zones over a year; see text.

3.2.5 Orientation of Focal Plane and Slits against the Sky

Spitzer has very limited ability to rotate the focal plane of the telescope. At any given time, the center of the sunshade (X-Z plane; see Figure 3.3) is always kept within ±2º of the Sun. This limited freedom in the "roll" angle is used to maintain a consistent orientation for long observations (AORs) and is not selectable by the observer. For each target, the orientation of the focal plane on the sky is a function of the position of the spacecraft along its solar orbit (i.e., the date of observation) and of the ecliptic latitude of the target. Targets in the ecliptic plane have only two possible focal plane orientation angles which are at 180º with respect to each other. Elsewhere in the sky, the focal plane has two ranges of orientation angles.

Science instrument aperture orientation within the focal plane is physically fixed. Because the spectrometer slits have quite different orientations within the focal plane (see Figure 2.1), this can be particularly limiting for IRS observations, but it also affects mapping with MIPS and IRAC.

3.2.5.1 Impact on Observers

Spitzer observations must be designed with this lack of roll control in mind. Observers who require a certain slit or array orientation angle relative to the target must check to determine when, if ever, the needed orientation is achieved, and must request a suitable range of observation date(s). It is possible that a particular target might not be observable at the same time of the year that a slit has the desired orientation. The focal plane/slit orientation on any date can be calculated and projected on the sky using the Spitzer observation planning tools (see Chapter 5).

Because constraining the orientation tightly also constrains the scheduling time tightly, it increases the risk that an observation may be subject to schedule conflicts. In addition, the presence of many highly constrained observations tends to reduce the observatory's efficiency. Therefore, observers should design observations that are as robust as possible to the maximum range of focal plane orientations, even if this makes the observations somewhat less efficient. The use of a timing constraint to obtain a specific orientation requires a *very* strong scientific justification when the observations are proposed. See the Call for Proposals for more information.

3.2.6 Bright Object Avoidance

No object outside the solar constraint zone poses a threat to instrument safety. However, an observer may wish to avoid observing the Earth and other bright moving objects to avoid compromising observations of faint targets. Therefore, the visibility windows calculated by Spot avoid certain bright moving targets by default, although the observer can choose to override the default.

With bright object avoidance turned on (the default), the visibility windows calculated for both inertial and moving targets will exclude regions of time when

the positions of (a) the Earth and Moon, and (b) a fixed list of bright moving objects (e.g., Jupiter, Saturn, bright asteroids) coincide with the target position. The visibility windows will be trimmed to delete any time periods when the Spitzer target is within 7º of the Earth or Moon and when the target is within 30′ of the other bright objects. The complete list of the bright moving objects is in Appendix 9.3.

The observer may choose to override the default (a) Earth/Moon or (b) other bright object avoidance. For example, to observe Jovian satellites, one would turn off (b) and leave (a) in effect.

The latest version of Spot includes the option to overlay all known moving objects on an image for a particular date. Observers concerned about bright moving objects not included in the list above may wish to use this option.

It should be noted that it is the observer's responsibility to check for any bright *inertial* targets in the field of view that might compromise the observation. There are some that will heavily saturate the instruments and therefore compromise not only that observation, but also possibly subsequent science observations as well. A list of some of the inertial objects known to saturate the instruments is available on the SSC website. You can also visualize bright objects near your AOR in Spot. Spot does *not* exclude observations of these objects. **The SSC reserves the right to exclude some targets and AORs as a result of impacts that bright objects would have on subsequent observations.**

3.3 Operational Constraints

Because Spitzer is formally committed to achieving certain operational goals (efficiency and mission lifetime), certain operational schemes must be adopted, which are described below.

3.3.1 Instrument Campaigns

Spitzer can only operate one science instrument at a time, and Spitzer observing is divided into dedicated instrument campaigns typically lasting from 7–21 days. Once the campaign plan for a particular observing cycle ("Baseline Instrument Campaign," or BIC) has been established, it is published on the SSC web pages. Minor changes in the instrument campaign plan in response to changes in the DSN schedule, on-orbit anomalies, ToOs, highly time-constrained observations, and changes to the SODB may be made. Updated BICs are posted regularly on the SSC website.

3.3.1.1 Length of Campaigns

There is a significant overhead associated with changeover from one instrument to another, due to the need to perform routine calibrations and engineering activities at the time of switch-off and switch-on. There are also performance limitations in some observing modes immediately following an instrument

changeover. For example, the IRAC electronics stabilize after a warm-up of ~1/2 hour after switch-on. In addition, there is a modest reduction in pointing accuracy following a change between IRS/MIPS (which share common warm electronics) and IRAC, due to thermal distortion when the heat load is shifted on the spacecraft.

In order to achieve the required mission observing efficiency and ensure good performance, changeovers will be made as infrequently as possible without sacrificing science needs. The best balance between efficiency and schedule flexibility is achieved by using each instrument for several (~7–21) days at a time. This generally permits long mapping projects, as well as highly time-constrained observations. The 7 to 21 day length is the expected typical case, not a rule. An instrument campaign may be interrupted in order to perform urgent ToO observations.

3.3.1.2 Sequence of Instruments

The instrument campaigns typically occur in the order IRAC-MIPS-IRS, which conserves helium and minimizes the effects on pointing performance described above. For Cycle-4, the instruments are operating in the following sequence: IRAC for ~10 days, MIPS for ~11 days, and IRS for ~14 days, with occasional deviations. Observations requesting MIPS followed immediately by IRAC are not feasible under this scheme. The length of these campaigns is determined by the relative fraction of observation requests in the SODB, and will probably change during the current and future Cycles.

3.3.1.3 Impact on Observers

Instrument campaign length and the order in which instruments are scheduled may ultimately be flexible guidelines, rather than unbreakable rules. However, the length and order will not normally significantly deviate from the scheme described above. If this changes in the future, campaign length and order will be carefully orchestrated to ensure maximum overall science return from Spitzer. Therefore, observations that necessitate unusually short campaigns and observations that necessitate interruptions to instrument campaigns impact the observatory efficiency and require *strong* justification. They may also incur substantial additional overheads, or even render the observations not schedulable. (See the Spitzer Observing Rules and the most recent Call for Proposals on the SSC website for details.)

ToOs, by their nature, can occur in the middle of the "wrong" instrument campaign; this is taken into account in the overhead assessed for ToOs that require rapid (~48 hour) response. The observer may want to take this into account when deciding how rapid a response to the proposed phenomenon is required to support the science goal of the observation. (See Spitzer Observing Rule #5.) **Rapid response ToOs that require a ~6 K telescope when the telescope is not already cold are unlikely to be accommodated.**

Other observations that strongly perturb the system are those that require sequential observations with all three instruments on a rapid timescale. If only two instruments are required, then it may be possible to schedule a small number of observations near a campaign boundary. If all three are required, then it will be necessary to switch-on, calibrate, observe, and switch-off in very short order. This is very inefficient and may also adversely impact the quality of the data obtained. This type of observation can be requested using the available scheduling constraints (see section 5.5.3), but requires a *very strong* scientific justification and will incur special overhead burdens.

3.4 Pointing Capabilities

3.4.1 Pointing Control System

The PCS includes the hardware and flight software necessary for precision telescope pointing, stabilization, slewing, tracking, and safe mode functions. The PCS performed the initial attitude acquisition of the spacecraft following launch vehicle separation. It provides periodic boresight calibration for the telescope. The PCS provides the capability for both rapid large angle slews and small maneuvers to place and reposition science targets within the science instrument apertures; it maintains the solar array orientation toward the Sun; and, it points the high gain antenna toward Earth for downlink. The PCS also contains Wide Angle Sun Sensors (WASSs) that act as a second check on the spacecraft orientation to ensure that the hard pointing constraints (see section 3.1) are not violated. High-level fault protection will place the telescope in a safe mode if a violation is detected. The performance numbers presented in this and the following sections are based on our present understanding and measurements of the on-orbit PCS performance.

The PCS is a celestial-inertial, three-axis stabilized control system. A high performance star tracker/inertial reference unit (ST/IRU) package provides attitude determination and reconstruction capabilities. On-board pointing commands and variables use the J2000 coordinate system. Reference to the J2000 celestial sphere is implemented within the ST through autonomous identification of stars carried in an on-board catalog of 87,000 Tycho stars down to 9th visual magnitude[12].

The ST is used to point an instrument boresight to a desired location on the sky with an initial accuracy of at least 0.5″ (1σ radial). The ST field of view is 5°×5°, which ensures that Spitzer can point to any part of the sky and have the ST meet its pointing requirements. Typically 40 stars are used simultaneously. The gyros provide pointing stability when not using the ST as a pointing reference; the pointing drift derived from the gyros is <3 mas/sec over 8 hours. The drift rate

---

[12] The on-board catalog actually uses the ICRS coordinate system, so, in fact, the on-board pointing system is really ICRS. However, the differences between ICRS and J2000 are so small (≤120 mas) that no conversions are made, and Spitzer is effectively considered to use the J2000 coordinate system.

when using the IRU-only mode should generally be better than 1 mas/sec over 200 sec.

All telescope pointing is defined and calibrated relative to redundant PCRSs located in the focal plane. During the course of the mission, the PCRS is periodically (about every 12 hours) used to calibrate the telescope-to-star tracker boresight alignment that may drift due to thermo-mechanical effects. Each PCRS detector is a Si PIN photodiode array divided into two 4×4 subarrays for redundancy. Each pixel is 250 μm square, with a plate scale of 10″ per pixel. The PCRS calibration measures the star position with an accuracy of 0.1″ (1σ per axis), and is sensitive down to 10th visual magnitude at a wavelength of 550 nm.

Spitzer also has WASSs, which measure the Sun's position with respect to the spacecraft. These sensors were used during initial attitude acquisition after launch, as well as for Sun avoidance, fault protection, and safe mode during the mission. Each wide-angle Sun sensor provides a field of view of 2π sr with an accuracy of ±0.1° at null. They are placed at the top and the bottom of the solar panels to maximize the coverage, with their boresights aligned to the spacecraft Z-axis.

Four reaction wheels provide the primary control actuation for all modes of operation. They are mounted in a pyramid orientation about the X-axis; each canted at 30° towards the X-axis. Over time, angular momentum accumulates in the reaction wheels, due primarily to the small offset between the center of mass and the center of (radiation) pressure. Unlike an observatory in low Earth orbit, which can dump this momentum magnetically, Spitzer has a "Reaction Control System" (RCS), which uses cold nitrogen gas thrusters to provide the reaction wheel momentum unloading capability; opportunities to dump momentum are scheduled during routine downlinks. The nitrogen supply is sized to accommodate a mission lifetime in excess of 5 years.

On-orbit measurements show that the PCS is capable of slewing the telescope 180° in 900 sec, 1° in 60 sec, and 1′ in 6 sec, while maintaining its inertial pointing knowledge. These times include the acceleration and deceleration of the telescope, but do not include the time it takes for the PCS to stabilize after the slew has completed. The pointing system has several operating modes, and the AOTs are designed to use the pointing mode most appropriate for each observing mode. Settling time varies with operating mode and slew magnitude. For IRU-only slews, slews less than 30′ settle to within 0.2″ rms within 10 sec. (Settling may take longer in some cases.) The AOTs make use of on-board slew completion and stabilization indicators to proceed with the observation as soon after a slew as is possible. Note that the time required for small slews, dithers, offsets, settling, etc., within an AOR is considered part of the observation. Time estimates provided through Spot account for these transparently and accurately.

The best available pointing system models are incorporated into Spot. However, note that Spot estimates may change if, e.g., Spitzer's slew rate changes.

3.4.2 Pointing Accuracy and Stability

The blind pointing accuracy is the same as the on-board attitude knowledge, <0.5″ (1σ radial rms) with a stability of <0.1″ (1σ radial rms over 200 sec). In the incremental pointing mode, the PCS performs controlled repositioning of the boresight with an offset accuracy usually no worse than 0.55″ (1σ radial rms) across angular distances of up to 30′ and usually better than 0.4″. This accuracy is sufficient to move a source from the IRS Peak-Up array to one of the spectrometer slits.

IRS observers can use the PCRS for acquisition of their science target; see the IRS chapter (chapter 7) for more information.

MIPS observers should note that the initial blind pointing accuracy for the 160 μm array is expected to be significantly worse than for the shorter wavelengths (~3.9″). See the MIPS chapter (chapter 8) for more information.

3.4.3 Scanning Stability and Performance

Spitzer can execute linear scans at selected rates from 0.01″/s to 20″/s. The MIPS Scan Map mode uses rates from 2″/s to 20″/s.

3.4.4 Tracking Capabilities

Spitzer does not have a true tracking capability for Solar System Objects (SSOs). However, Spitzer simulates tracking by scanning in linear track segments at rates up to 1″/s. The linear track segments are linear in equatorial coordinate space; they are commanded as a vector rate in J2000 coordinates, passing through a specified RA and Dec at a specified time.

The SSO ephemerides are maintained on the ground, rather than on-board Spitzer, but an observer may specify flexible scheduling constraints, and the linear pseudo-track specification (start point, rate and direction, in equatorial J2000 coordinates) is calculated at the time of scheduling. The observation is executed at the scheduled time (within a window of +3, -0 s), and the PCS follows the track as specified, assuming the given start point corresponds to the time given in the track command. All other PCS movements can be superposed on a specified track, including dithers and scans.

PCS measurements indicate that the blind track acquisition accuracy is ≤ 0.5″, independent of the track rate, which is consistent with expected performance on fixed targets. The track stability is better than scan mode requirements: ~0.5″ in 1000s. The offset accuracy during tracking is better than ~0.55″, also consistent with the required performance on fixed targets.

Section 5.6 presents more detailed information on observing Solar System Objects, including ephemeris management.

3.4.5 Pointing Reconstruction

Pointing reconstruction refers to the post-facto determination of where the telescope was pointed with a greater accuracy than was known by the flight system. Spitzer has a requirement to perform pointing reconstruction to $1.4''$. Based on in-orbit results, Spitzer meets this requirement, except for observations scheduled immediately after an instrument changeover (in which case the inaccuracy may be as much as $\sim2''$). The nominal pointing is reported in the data products. Based on post-BCD processing and comparison of extracted sources with 2MASS sources, users will receive data with IRAC data pointing reconstruction generally $<1''$, and MIPS scan pointing reconstruction generally $<1.5''$, with respect to the 2MASS coordinate system. Relative pointing (relative separation of two objects within an AOR) is good to $<0.5''$.

4 Telescope and Focal Plane

4.1 Optical Design

The telescope is a Ritchey-Chrétien design, known from on-orbit measurements to be diffraction-limited at 5.5 μm over the IRAC field of view, and expected to be diffraction-limited at 6.5 μm over the entire field of view. The 85-cm diameter primary mirror and the rest of the telescope structure are fabricated entirely of beryllium, utilizing advances in optical design, testing, and fabrication to produce a lightweight telescope that operates at cryogenic temperatures. The optical design parameters for the telescope are summarized in Table 4.1 below. The telescope configuration is shown in Figure 4.1.

Note that this chapter addresses the telescope itself; the point spread function (PSF) at each of the science instruments is discussed in the corresponding instrument-specific chapter (chapter 6, 7 or 8).

Table 4.1: Telescope Configuration

Optimal Parameter Description	Value at 5.5 K
System Parameters:	
Focal Length	10,200 mm
Focal Ratio	f/12
Back focal length (PM vertex to focus)	437 mm
Field of View (diameter)	32.0 arcmin
Field curvature radius	140.5 mm
Wavelength coverage	3 μm–180 μm
Aperture Stop:	
Location	Edge of primary mirror
Diameter of OD obscuration	850.00 mm
Diameter of ID obscuration	320.00 mm
Linear obscuration ratio	0.3765
Primary Mirror (hyperbola)	
Radius (concave)	-2040.00 mm
Conic constant	-1.003546
Clear aperture	850.00 mm
Focal ratio	f/1.2
Secondary Mirror (hyperbola)	
Radius (convex)	-294.343
Conic constant	-1.531149
Clear aperture (OD)	120.00 mm
PM to SM spacing (vertex to vertex)	887.528 mm
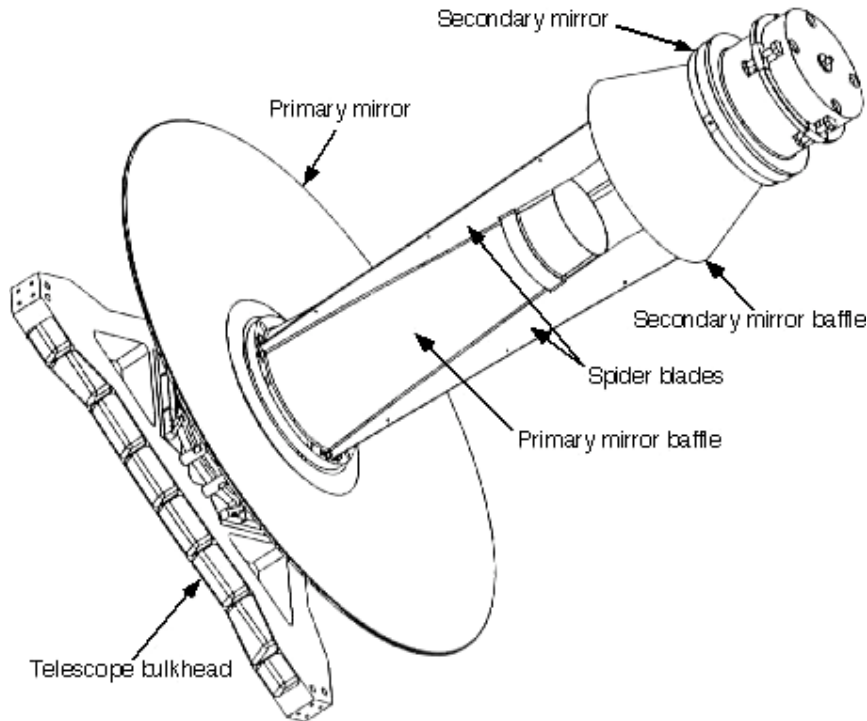
Figure 4.1: Spitzer Telescope Assembly.

The telescope employs a single arch primary mirror, which reduces mass. The primary mirror is supported on three bipod flexures relatively close to its axis. The spider blades, primary mirror baffle, and secondary mirror baffle are integrated into a one-piece, relatively small diameter metering tower extending through the central hole in the primary mirror.

4.2 Optical and Thermal Performance

4.2.1 Surface Accuracy

The Ritchey-Chrétien design minimizes spherical aberration and coma over large fields of view. Field curvature varies quadratically with field angle. Similarly, the rms wavefront error at best focus varies quadratically with field angle and equals 0.52 waves ($\lambda = 0.6328$ μm) at the edge of the field. Essentially all of this error is due to astigmatism.

The surface figure for the primary mirror was measured at cryogenic temperatures to be 0.067 μm rms over the entire clear aperture, meeting the specification of 0.075 μm rms (Figure 4.2).
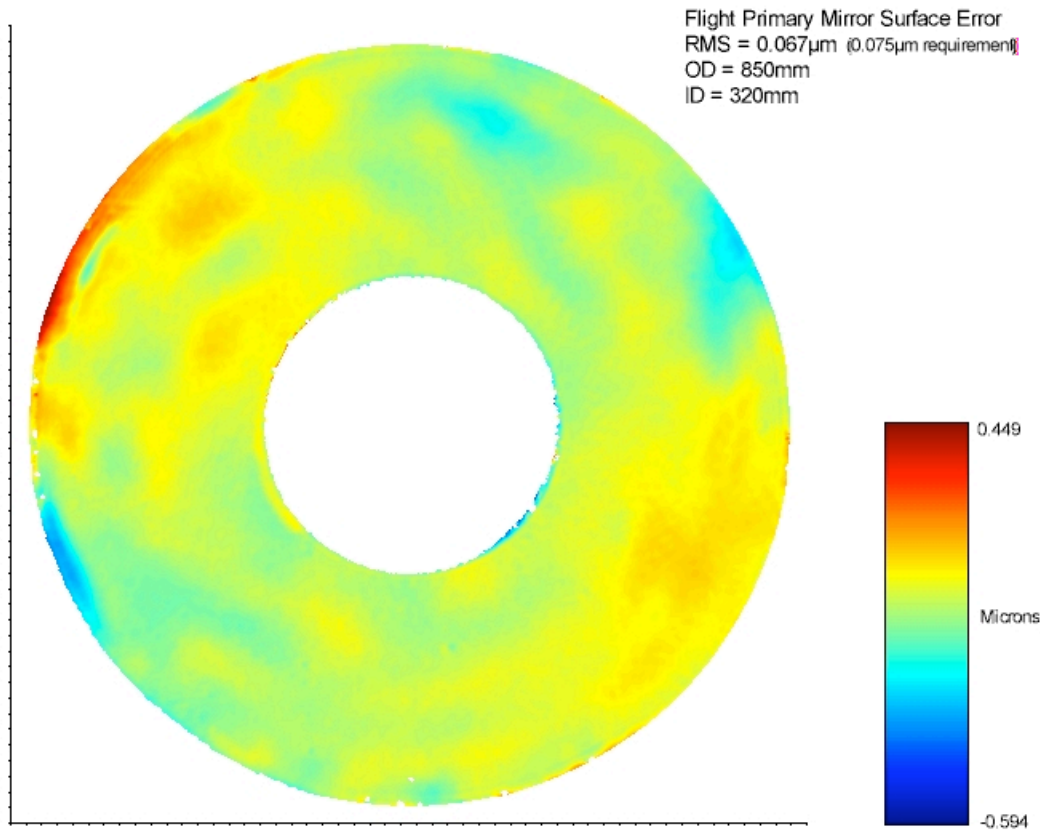
Figure 4.2: The deviations in the flight primary mirror surface. The RMS error was measured at cryogenic temperatures to be 0.067 µm rms over the entire clear aperture, meeting the specification of 0.075 µm rms.

4.2.2 Wave-front Errors

The telescope is required to provide a beam to the telescope focal surface that is diffraction limited (transmitted wave front error $< \lambda /14$ rms) at 5.5 µm over the entire field at operating temperature. At a wavelength of 3.5 µm, the telescope produces a wave front error of less than 0.13λ rms over the IRAC field of view, and the image of a point source contains 45% of the encircled energy within a diameter of 2″. The actual performance of the telescope is slightly better than the requirements; see IRAC chapter (specifically, Table 6.1) for more information.

4.2.3 Throughput

The requirements for the telescope assembly are that it shall provide a minimum end-of-life throughput no less than that given in Table 4.2. Telescope throughput is defined as the ratio of energy from a point source reaching the telescope focal surface to the energy collected by an 85 cm diameter mirror. Factors that degrade telescope throughput are the central obscuration (including spiders), mirror reflectivity as a function of wavelength and losses due to contamination.

Table 4.2: Telescope Throughput

Wavelength (μm)	End of Life Throughput
0.55	> 0.18
3.5–6.5	> 0.70
6.5–10	> 0.73
10–20	> 0.74
20–200	> 0.75

4.2.4 Stray Light Rejection

The cryostat, telescope, multiple instrument chamber and science instruments are designed and baffled such that, at all wavelengths from 3.6 μm to 160 μm (center of passbands), celestial stray radiation and internal stray radiation:

- Do not, except for lines of sight near bright sources, increase by more than 10% the photon noise of the natural background in the direction of the line of sight of the telescope. This requirement implies that the combination of celestial stray radiation and internal stray radiation must be <21% of the natural background at the instrument detector arrays.
- Display no gradients or glints in the celestial stray light that will increase confusion noise over natural levels or produce false sources.
- Do not significantly decrease the contrast of the first dark ring of the diffraction-limited point spread function.

The conformance of the Spitzer design to its stray light requirements was verified by analysis using the APART stray light analysis program and an analytical test source designed to approximate the brightest celestial source expected in each of Spitzer's wavelength bands. The actual scattered light performance of the Spitzer telescope was characterized on-orbit during the early parts of the mission and will continue to be characterized during nominal operation of the telescope. In some cases, there are modifications or caveats to the requirements above. Discussion of stray light as it pertains to each instrument is covered in the instrument chapters later in this document.

A useful output from the stray light analysis is a set of predicted point source transmission curves. The point-source transmission function (PST) is the inverse of the ratio of the flux density (W/m^2/Hz) of an off-axis source to the flux density at the telescope focal plane due to light scattered from that source. The separate PST curves in Figure 4.3 refer to different azimuthal locations of the celestial point source. An azimuth of 0º refers to the anti-Sun direction. The differences among the azimuths are mainly due to the changing illumination of struts supporting the secondary mirror. Figure 4.4 shows the variation as a function of wavelength.
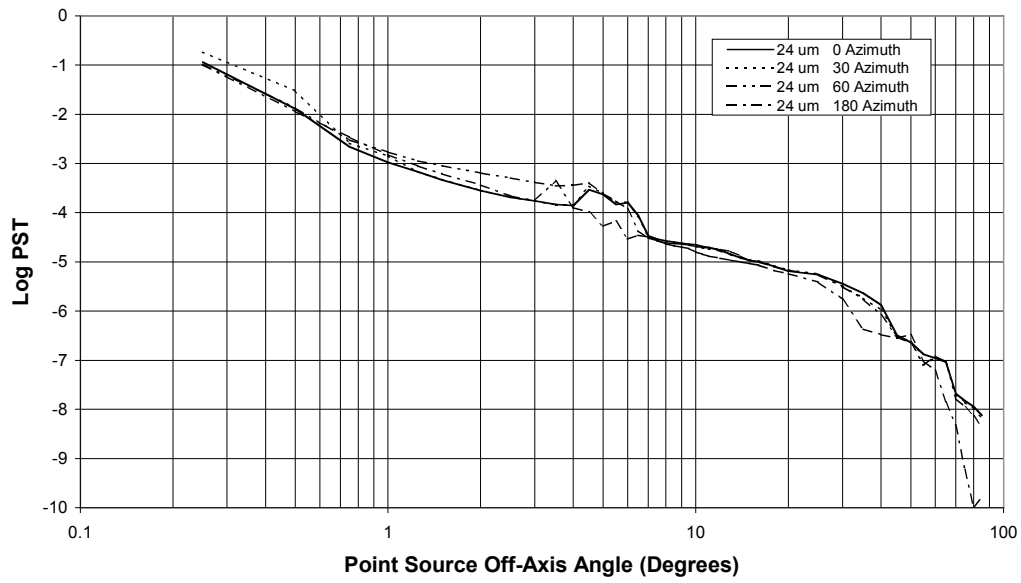
Figure 4.3: Theoretical PST for 24 μm off-axis as a function of azimuth.
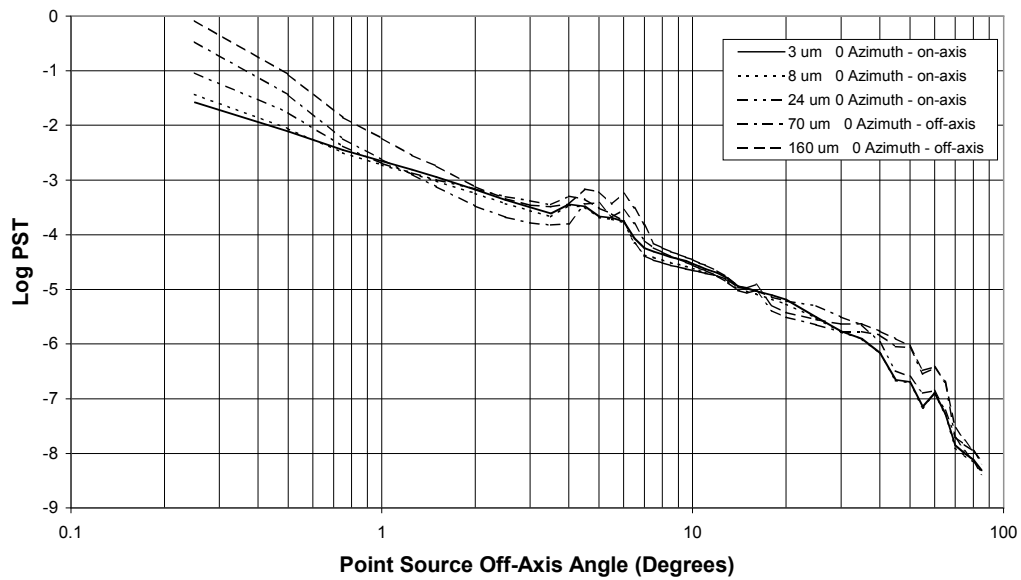
Figure 4.4: Theoretical PST as function of wavelength for fixed azimuth on-axis.

These plots can be used to estimate the stray light contribution from a given source. For example, at 8 µm, Vega has a flux density of ~62 Jy. From Figure 4.4, the 8 µm PST at 1º off axis is ~2×10^{-3}, so the predicted flux density in the Spitzer focal plane due to Vega at 1º off axis is 120 mJy[13].

4.2.5 Telescope Temperature/Thermal Background

The telescope is cooled by helium vapor vented from the cryostat. The helium vaporization is driven by power dissipated by the Science Instrument cold assemblies. A supplemental heater is available to provide additional vaporization, if needed. The required telescope temperature and heat load drive the required helium flow rate. As discussed in section 3.1.1, the telescope was launched warm and gradually cooled by a combination of radiative cooling and helium boil-off to ≤ ~6 K during the first ~45 days in orbit. The telescope temperature varies between <~6 K and 12 K depending upon how much power is dissipated in the cryostat. The power thus dissipated depends primarily upon the science instrument in use; the "make-up" heater may be used to ensure that the telescope is maintained at ~6 K when needed for 160 µm observations.

4.2.6 Telescope Focus

Spitzer is equipped with a secondary mirror focus mechanism, which was operated both on the ground and on orbit. Prior to launch, the end-to-end image quality was measured on the ground and the mechanism was set to the position that was predicted to give optimal focus following in-orbit cool-down. During In-Orbit Checkout (IOC), it was found that the telescope focus point was about 1.85 mm above (toward the back of the primary mirror) the optimum focus position for the science instruments. Thirty-eight days after launch, when the telescope had become thermally stable, the secondary mirror was moved toward the primary mirror to bring the telescope into focus. The instruments are confocal, so that separate adjustment of the focus for each instrument will not be necessary during the remainder of the Spitzer mission.

4.3 Focal Plane Layout

Figure 4.5 shows the actual location of the pick-off mirrors that feed the science instruments as viewed from above (looking down the boresight). The Y- and Z-axis directions are the same as the telescope coordinate system described in Chapter 3; the +X direction comes out of the page. Figure 2.1 shows the Spitzer entrance apertures as projected onto the sky and appears inverted compared to Figure 4.5, due to the combined effects of looking out from behind the focal plane and the projection of the sky onto the focal plane through the telescope optics. (To understand this inversion, recall the 3-D geometry and the fact that a Cassegrain telescope inverts its image.)

---

[13] The natural background at 8 µm near an ecliptic pole is 5.3 MJy/sr, which is imaged to a flux density of 5.3 MJy/sr × π /(4×12^2) = 2.9×10^4 Jy in the focal plane of the *f*/12 Spitzer telescope. The scattered light from Vega at 1º off-axis is far below the natural background. The first stray light trouble from Vega at 8 µm should come from the outer parts of its diffraction-limited PSF.

Nominal Aperture or Pick-Off Mirror
Locations Viewed Down the Telescope
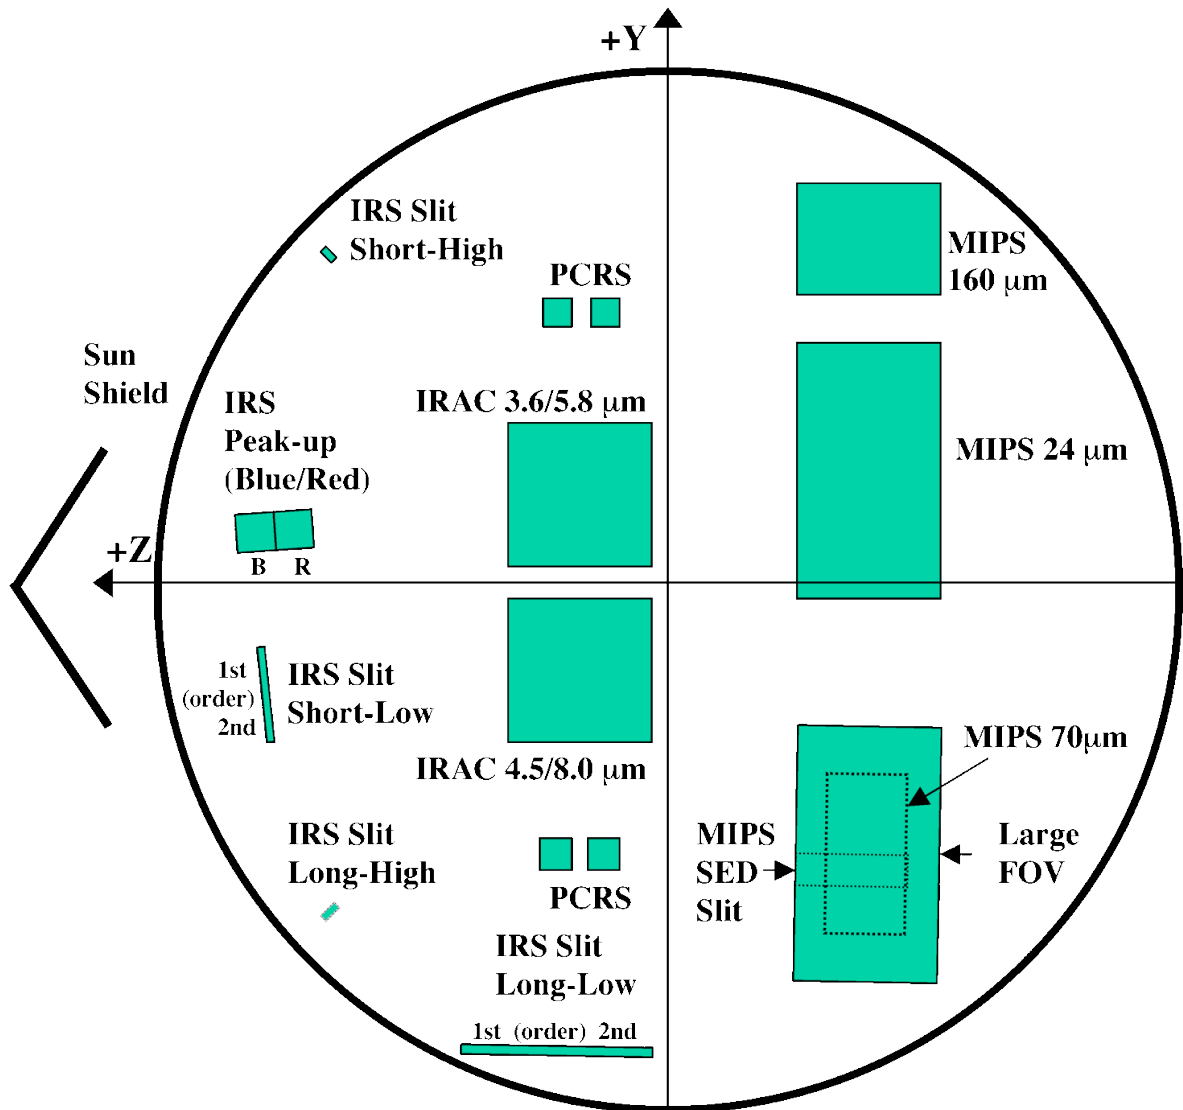
Nominal Field Radius = 16 arcminutes

Note: This is an approximate mechanical layout of the focal plane as viewed from above. It is not a projection onto the sky.

Figure 4.5: Schematic view of Spitzer focal plane from above, looking down the boresight. The solar panel is on the IRS side of the spacecraft. This figure shows the region of the focal surface where the pick-off mirrors for each instrument are located. This is in contrast to where the apertures project onto the sky. See Figure 2.1 for comparison.

Table 4.3: Spitzer focal plane layout: approximate offsets from boresight to aperture centers

Aperture	Z (′)	Y (′)
IRAC 3.6 μm	2.86	3.67
IRAC 3.6 μm subarray	4.97	5.88
IRAC 5.8 μm	2.86	3.61
IRAC 5.8 μm subarray	0.72	5.71
IRAC 4.5 μm	2.72	-3.05
IRAC 4.5 μm subarray	4.80	-0.92
IRAC 8.0 μm	2.75	-3.07
IRAC 8.0 μm subarray	0.65	-0.97
IRS SL 1st order	12.03	-2.79
IRS SL 2nd order	11.91	-4.09
IRS LL 1st order	4.41	-14.00
IRS LL 2nd order	1.21	-14.07
IRS SH	10.56	10.01
IRS LH	10.42	-10.23
IRS Red Peak-Up	11.63	2.00
IRS Blue Peak-Up	13.52	1.91
MIPS 24 μm	-6.72	4.25
MIPS 70 μm default	-6.53	-8.06
MIPS 70 μm fine	-7.10	-6.81
MIPS 70 μm default side A	-7.95	-8.11
MIPS 70 μm fine side A	-6.39	-6.71
MIPS 160 μm	-7.35	12.11
MIPS SED	-6.06	-9.44

Table 4.3 gives the measured offsets from the center of the field of view to the centers of the main Spitzer science apertures. The values were determined post-launch based on the results from the focal plane mapping survey during IOC+SV.

Each of the fields of view that are used by Spitzer has a code that specifies the position in the focal plane. The uplink software uses the field-of-view index for commanding the target position, and the pipelines similarly use it to reconstruct where the telescope was pointing. These indices and field-of-view names appear in the headers of Spitzer data, and are listed in Table 4.4.

Table 4.4: Field-of-View Indices

FOV name	FOV index
Telescope Boresight	1
PCRS 1A	4
PCRS 1B	5
PCRS 2A	8
PCRS 2B	9
IRS red peak-up FOV center	18
IRS red peak-up FOV sweet spot	19
IRS blue peak-up FOV center	22
IRS blue peak-up FOV sweet spot	23
IRS SL 1^{st} order 1^{st} position	26
IRS SL 1^{st} order 2^{nd} position	27
IRS SL 1sr order center position	28
IRS SL module center	29
IRS SL 2^{nd} order 1^{st} position	32
IRS SL 2^{nd} order 2^{nd} position	33
IRS SL 2^{nd} order center position	34
IRS LL 1^{st} order 1^{st} position	38
IRS LL 1^{st} order 2^{nd} position	39
IRS LL 1^{st} order center position	40
IRS LL module center	41
IRS LL 2^{nd} order 1^{st} position	44
IRS LL 2^{nd} order 2^{nd} position	45
IRS LL 2^{nd} order center position	46
IRS SH 1^{st} position	50
IRS SH 2^{nd} position	51
IRS SH center position	52
IRS LH 1^{st} position	56
IRS LH 2^{nd} position	57
IRS LH center position	58
IRAC center of 3.6 and 5.8 arrays	67
IRAC center of 3.6 array	68
IRAC center of 5.8 array	69
IRAC center of 3.6 subarray	70
IRAC center of 5.8 subarray	72
IRAC center of 4.5 and 8.0 arrays	74
IRAC center of 4.5 array	75
IRAC center of 8.0 array	76
IRAC center of 4.5 subarray	77
IRAC center of 8.0 subarray	79
IRAC between arrays	81
MIPS 160 center	87
MIPS 160 large only	89
MIPS 160 small	91, 92

FOV name	FOV index
MIPS 24 center	95
MIPS 24 small	99, 100
MIPS 24 large	103, 104
MIPS SED	105, 106, 121, 122, 123, 125, 126
MIPS 70 center	107
MIPS 70 TP	109
MIPS 70 default small	111, 112
MIPS 70 scan	113
MIPS 160 scan	114
MIPS 70 default large	115, 116
MIPS 70 fine	117, 118, 119, 120, 124, 127

5 Observing with Spitzer

The Spitzer observer's interface to the observatory, including the instruments, is the set of eight AOTs. Spitzer's science instruments are relatively simple in the sense of having few modes and even fewer moving parts (only two — the IRAC shutter, which is not planned to be used in regular operations, and the MIPS scan mirror). AOTs are a central design concept in Spitzer's science operations. The use of relatively simple parameterized observing modes enhances the reliability of Spitzer observations and calibration, improves the archival value of Spitzer data, and reduces cost.

This chapter describes the Spitzer AOTs, introduces the concept of an AOR, introduces the SSC-provided tools for planning and creating observations, and describes some of the rules and procedures related to observing with Spitzer. Special types of observations, such as Solar System Objects (SSOs), ToOs, and Generic Target observations are also discussed.

The tool which Spitzer observers must use for planning observations and submitting proposals is Spot, which is described in section 5.3.

5.1 The Eight Spitzer Astronomical Observation Templates – AOTs

An AOT is a distinct Spitzer observing mode; there are eight Spitzer AOTs for the three science instruments.   Five of the AOTs (IRAC Mapping, IRS Staring, IRS Spectral Mapping, MIPS Photometry, and MIPS Scan) were commissioned during the first three months following Spitzer's launch. Another AOT (MIPS Spectral Energy Distribution [SED] mode) was available for the Cycle-1 Call for Proposals, though observations were not scheduled until the AOT was commissioned in October 2004. Two AOTs became available, beginning in Cycle-2: MIPS Total Power [TP] mode, and IRS Peak-Up Imaging (PUI). All of these existing AOTs are available for Cycle-5.

The observing modes are:

InfraRed Array Camera (IRAC) Mapping/Photometry
    The IRAC AOT is used for simultaneous imaging at 3.6, 4.5, 5.8 and 8.0 μm, over the two ~5.2′ × ~5.2′ fields of view.
InfraRed Spectrograph (IRS) Staring-Mode Spectroscopy
    The IRS staring mode is used for low-resolution long-slit spectroscopy (R= 64–128) from 5.2 to 38.0 μm and high-resolution spectroscopy (R~600) from 9.9 to 37.2 μm. The IRS Staring mode also supports raster mapping. It also returns images from the IRS Peak-Up array, which has a field-of-view of approximately 1′ square and two filters covering 13.5–18.5 μm and 18.5–26 μm.
IRS Spectral Mapping
    The IRS Spectral Mapping AOT is used to perform slit scanning spectroscopy for fields up to a few arcminutes in extent.

IRS Peak-Up Imaging

The IRS Peak-Up Imaging AOT provides imaging only using the Peak-Up array, which has a field-of-view of approximately 1′ square and two filters covering 13.5–18.5 μm and 18.5–26 μm.

Multiband Imaging Photometer for Spitzer (MIPS) Photometry and Super Resolution Imaging

The MIPS Photometry and Super Resolution AOT is used for imaging photometry and high-resolution imaging at 24, 70 and 160 μm. An "enhanced" small-field mode for 160 μm observations is also available.

MIPS Scan Mapping

The MIPS Scan Map AOT is used for large field maps at 24, 70 and 160 μm. The maps are constructed using slow telescope scanning, combined with motion compensation using a cryogenic scan mirror. Maps are built up of ~5′ (2.5′ for full coverage at 70 μm) wide strips between 0.5° and 6° in length.

MIPS Spectral Energy Distribution (SED)

The MIPS Spectral Energy Distribution AOT is used for very low-resolution (R=15–25) spectroscopy covering 52-97 μm using the MIPS 70 μm Ge:Ga array.

MIPS Total Power Measurement (TP)

The MIPS Total Power Mode AOT provides zero-level-reference observations for absolute brightness of extended sources.

Each AOT and its usage are discussed in detail within the respective instrument-specific chapter. Examples of how to create observations using the AOTs are also included in the instrument-specific chapters 6 through 8 and are more extensively discussed in the Spitzer Observation Planning Cookbook.

5.2 Astronomical Observation Request – AOR

When all the relevant parameters for an AOT are specified and linked to a description of the target, the resulting fully specified observation is called an AOR, which is the fundamental unit of Spitzer observing. AORs are created using Spot (see section 5.3). An AOR can be thought of as a list of parameters that, when properly interpreted, completely describe an observation. In fact, an AOR, as represented in the SODB, contains a series of keywords and values that are used to create the sequence of commands that are sent to the observatory to carry out the observation. This ASCII format is what Spot actually creates and stores in its output file. The parameters of the AOT are specified in Spot by filling out a form, sometimes called an "AOT Front End."

The AOR is the fundamental scheduling unit for Spitzer: An AOR cannot be subdivided, will not be interrupted for other activities (such as downlinks), and is handled as a unit by the observatory. Because of this non-interruptible nature and the need to perform certain activities periodically (e.g., detector anneals, pointing system calibrations, and downlinks) a maximum duration exists for an AOR. That maximum time is 3 hours for MIPS AORs and 8 hours for both IRAC

and IRS AORs. Longer observations can be specified using multiple AORs and relational constraints to identify these AORs as members of a related group (see section 5.5.3). When specifying such related groups, it should be taken into account that, in general, several AORs of the maximum length cannot be placed back-to-back. Even taking into account only that the spacecraft must downlink every 12 hours and perform pointing calibrations every 8 hours (ignoring calibrations and maintenance of the science instruments), it can be seen that the observatory schedule may be broken up into irregular chunks of time. Observing programs that are composed of short, independent AORs that can be scheduled any time their targets are visible are easiest to schedule. Observing programs that require building related groups of constrained observations are more likely to fit into a schedule if the AORs themselves are short and the constraint is constructed to allow maximum flexibility (for details on constraints see section 5.5).

An AOR contains three categories of information:

Astronomical Target
> The target of an AOR can be a single pointing or a cluster of pointings within a 1° radius, at which the specified observation is repeated identically. The single pointing or cluster may be either of an inertial target or a moving target.

AOT-Specific Parameters
> As the name implies, these vary from AOT to AOT. They include instrument configuration, exposure time and dedicated mapping parameters.

Timing and Relational Constraints
> These constraints represent scheduling directives for an AOR or for a related group of AORs. The details of the kind of constraints that are supported are in section 5.5.3. Timing constraints are used to specify a window when an AOR should be executed (e.g., to observe a comet at maximum solar elongation). Relational constraints are used to specify how AORs within a group are related to one another (e.g., a series of AORs that define a very deep map and must be executed consecutively).

5.3 Science User Tools

Science User Tools are software packages and other materials (such as tables and graphs) that are provided by SSC to help the astronomical community plan, prepare, submit, monitor, and interpret the results of their Spitzer observations. They are all available from the SSC website, specifically the Proposal Kit area (http://ssc.spitzer.caltech.edu/propkit/). We highlight a few tools here for illustrative purposes. Note: there is also a new perl-based tool, available from the Proposal Kit area, allows proposers to format the Observations Summary Table portion of the proposal template, using the AORs to be submitted.

5.3.1 Spot

Spot is a multi-platform JAVA-based client-server GUI-driven software tool intended to assist potential and approved Spitzer observers in planning and modifying their observations. Spot, along with the Spot User's Guide, can be

downloaded from the Proposal Kit section of the SSC web pages
(http://ssc.spitzer.caltech.edu/propkit). It allows investigators to construct and
edit detailed AORs by entering targets and selecting from a variety of preset
instrument-specific functions (e.g., exposure times, instrument modes, dither
patterns, and observing constraints). Spot also includes useful visualization tools
to permit an investigator to see how proposed Spitzer observations and the Spitzer
focal plane is laid out on the celestial sky. These capabilities allow observers to
retrieve relevant images from other astronomical surveys (in any of a number of
wavelengths) and archives. It calculates estimates of Spitzer observing time
(including telescope overheads) for each AOR in a proposed program, along with
target visibility information, focal plane position angle for a selected observation
date, and estimates of the zodiacal and cosmic infrared background at the target.
Spot allows investigators to view AORs from all previously approved programs;
it also supports "check-in" of active Spitzer programs for approved modifications
by investigators (see section 5.4).

Spot allows prospective investigators to plan, develop and modify their proposal
in an iterative manner. That is, a proposer can write a portion of their proposal
and define their accompanying AORs, save the results locally, and then re-load
those results at a later time for subsequent modification. The saved cover sheet,
proposal text, and AORs can be modified repeatedly before the proposal
submission deadline.

*Prospective investigators are encouraged to download Spot and to start planning
their Spitzer observing programs well before the proposal submission deadline.*

Spot is currently supported on multiple Windows, Linux, Sun Solaris, and
Macintosh platforms. For details, see the Spot User's Guide and Release Notes,
available at the SSC website with the downloadable software. While running
Spot, a network connection to the Spitzer Science Center is required to obtain
observing time estimates, visibility, orientation, or background estimate
information, and to submit and update proposals; a network connection is not
necessary for selecting observation parameters.

5.3.2 IRSky & IBIS

Although Spot is recommended for planning Spitzer observations, the stand-alone
tool IRSky (or its batch version, IBIS) is also available for determining
background estimates. For some observations (e.g. near the Galactic plane; see
below), IRSky produces better background estimates than Spot.

IRSky is an X-Windows based package accessible through the SSC Proposal Kit
web page. An X-Windows emulator is required to use IRSky from a PC.

Spot provides overlays of coverage maps for Spitzer, whereas IRSky does not.
As in Spot, IRSky will give the observer an indication of how the Spitzer focal
plane will appear on the IRAS Sky Survey Atlas (ISSA) 12, 25, 60, and 100 μm
maps of the sky for specific observations. However, this task is easier in Spot,

since the necessary input position angle can be obtained for a particular target and date. This is not possible for IRSky.

Note that IRSky estimates the background emission using a different model from that used by Spot. The Spitzer background estimator used in Spot is based on the (pre-launch) memo from 2000, "SIRTF Background Estimation: Methods and Implementation," by W.T. Reach, available on the SSC website. The IRSky background estimator uses a zodiacal background model and an interstellar cirrus background estimator derived from IRAS data. The model incorporated into Spot is based on both IRAS and COBE data, takes into account the time of observation in calculating the zodiacal light background component, and uses improved wavelength-interpolation schemes. However, near the Galactic plane or other very bright extended emission, the IRSky estimator is expected to give better results than the background model in Spot. Background estimates are an important part of a Spitzer proposal and can be complicated. A collection of links relevant to Spitzer background estimation can be found as part of the Infrared Compendium on the SSC website.

5.3.3 Performance Estimation Tool (PET)

The most important web-based SSC tool is the Performance Estimation Tool (PET). The PET is actually a set of Javascript tools to aid in designing Spitzer observations. There are four PET components: The Sensitivity PET, the Spectroscopy PET, the Extragalactic PET, and the Stellar PET. *Users of the PET should pay close attention to the various notes and warnings on both the main PET pages and the help pages*.

The Sensitivity PET (SENS-PET) is an imaging sensitivity estimator. It takes as input IRAC, IRS PUI, and/or MIPS imaging instrument configurations, and a background level. It produces as output the instrument sensitivities (both for point source and extended objects), and the total exposure depth per pixel.

The Spectroscopy PET (SPEC-PET) is designed specifically for spectroscopic sensitivity estimates. Users configure IRS and/or MIPS SED observing parameters, and the SPEC-PET returns an estimate of the instrument sensitivity.

The Extragalactic PET (EX-PET) makes predictions for imaging of extragalactic sources. As input, users may choose an SED model, background level, and IRAC + MIPS instrument configurations. The output includes flux in the instrument passbands, instrument sensitivities, S/N, and total exposure depth per pixel.

Finally, the Stellar PET (STAR-PET) predicts fluxes for imaging of stars. For input, users select a stellar spectral type + MK class, K-band magnitude or K-band flux density. The tool calculates the expected flux density in the IRAC 3.6, 4.5, 5.8, 8.0 μm passbands, and at 15 μm (IRS peak-up blue channel and short wavelength module), and 24 μm (IRS peak-up red channel and long wavelength module, and MIPS Si:As array).

5.3.4 Flux Unit Converter

Although Spitzer predictions and products are in Janskys, some astronomers may be more familiar with magnitudes. Another Javascript tool available at the SSC website converts back and forth between the flux (density) units of magnitudes and Janskys, as well as a variety of other flux (density) units. See also the Appendix of this document for more information.

5.3.5 Leopard – Spitzer Archive Interface

The observer's interface to the archive is called Leopard, and is very similar to Spot (see above). Leopard comes with the Spot software download on the Proposal Kit web pages. Detailed information about how to retrieve public and/or proprietary data using Leopard can be found in the Leopard User's Manual, also available on the Proposal Kit web pages. Observers should also use Leopard as their primary source for checking the Reserved Observations Catalog.

5.4 Modification of Observing Programs

The SSC envisions that it will execute most observations exactly as they are specified in response to the Call for Proposals. Therefore, the SSC *will expect observers to provide sufficient information at the time of proposal submission to implement the observations* (i.e., completed AORs); also, see the current Call for Proposals for more information.

In cases of conflicts or special needs, it will be possible for observers to request a modification of their (not yet scheduled) AORs via the Spitzer Helpdesk. *The request must be accompanied by a strong justification for the proposed changes.* Major modifications by GO observers are only allowed in exceptional circumstances and only outside of blackout periods. Minor modifications can be made at any time. See the SSC web pages for details on blackout periods and the definition of major vs. minor modifications. An observer can, at any time, request (again with strong justification) that a scheduling "hold" be placed on observations that need to be delayed. Only certain types of modifications are permitted, and the observer should check the Spitzer Observing Rules prior to requesting a modification of his or her observing program.

Upon completion of approved modifications, the observer will need to resubmit the revised AORs. Viewing, modification, and resubmission of AORs can all be done within Spot. Modification of programs requires the program username and password set at the time of proposal submission, once Science User Support has reactivated the password. The Science User Support Team at the SSC will review the changes and re-check for duplications prior to re-releasing the AORs for scheduling. Once a program has been resubmitted, the access for modifications will be disabled, and any further modification will require a new request for access, accompanied by supporting justification. Check the Science User Support section of the SSC web pages for full details on program modification.

5.5 Scheduling Considerations

The technical issues of telecommunications and telescope performance that drive science instrument campaign planning and the basic scheduling timescales are addressed in Chapter 3. In this section, the planned scheduling process is addressed from the observer's point of view.

The use of scheduling constraints is also discussed in this section. Scheduling constraints restrict the ability of the planning and scheduling system to create efficient schedules and to ensure the feasibility of scheduling all AORs. Therefore, a scientific justification is required for the use of scheduling constraints and must be provided in the relevant science proposal at the time of submission.

5.5.1 Scheduling Methodology

The Spitzer scheduling process constructs short-term observing schedules of one week in length. Each week's observing schedule is finalized about five weeks in advance of execution, to allow for processing of the command sequences by the Mission Operations team at JPL and review of the resulting command products. Scheduling is done one week at a time, regardless of the instrument campaign schedule, and a one-week schedule generally contains observations using more than one science instrument. The one-week schedule loads are uplinked weekly, following final review and project approval. In general, Spitzer "weeks" run from Thursday to Wednesday. After a schedule has been finalized, changes are only made in unusual circumstances (e.g., activation of ToO programs). Due to the fact that Spitzer is in contact with the ground for only a brief period every 12–24 hours, it is not practical to make changes to the on-board schedule on very short time-scales; the fastest a ToO can be accommodated is about 48 hours from the time of SSC's receipt of the approved AORs.

Observations are grouped into instrument campaigns, which are periods of several days during which only one instrument is used. Instrument campaigns will vary in length as needed to accommodate the science program, but it is expected that typically they will last about 5 to 21 days. In Cycle-4, IRAC campaigns are ~10 days, followed by ~11-day MIPS campaigns, and finally by ~14-day IRS campaigns (see section 3.3.1).

Rules regarding interruption of instrument campaigns and schedule changes on short notice are discussed in more detail in the Spitzer Observing Rules. Observers should also consult the Observing Rules and the SSC website for information on proposing programs seeking frequent instrument changeovers, as scheduling such programs is difficult or even impossible, given the operational constraints on the spacecraft (see section 3.3). Observers considering such programs should also be aware of hardware-related performance issues pertaining to instrument changeovers (see section 3.3).

Proposers for Cycle-5 should also see the Cycle-4 and Cycle-5 Scheduling memo available on the Proposal Kit webpage.

The planning and scheduling process includes determining the optimum structure of instrument campaigns. This type of campaign planning will occur shortly after a new set of observations becomes available following solicitation, review, and selection (i.e., a Call for Proposals). The BIC is published on the SSC website once it is considered stable, but it is subject to change without notice, due to the need to respond to actual events on-orbit.

5.5.2 Notification That an AOR Has Been Scheduled

Approximately three weeks before a schedule has been uplinked to the spacecraft, the list of scheduled AORs and their nominal execution times will be published on the SSC website. Unless an absolute scheduling time has been specified at the time of scheduling, the actual execution time of an AOR may differ from the scheduled time by a small amount (possibly up to several minutes). This is due to the non-deterministic nature of the on-board indicators for slew completion and settle time.

The appearance of a particular AOR in the published schedule is not the only notification to observers that the AOR has been officially scheduled. An email is sent to the observer notifying him/her that his/her observation has been scheduled. Once an AOR has been scheduled, only a significant anomaly (e.g., a missed downlink opportunity or a safing event) or a rapid-turnaround ToO could cause the schedule to change. Observers will generally not be notified if this happens. **Only under extreme circumstances can AORs be modified after they have been scheduled, but this is *highly* discouraged.**

In the case of AORs that are highly time-constrained for science reasons, if a specific sequence or dates have been provided to the observatory in advance, the Observatory Planning & Scheduling Team (OPST) will make every effort to schedule those AORs accordingly. This information should be provided at the time of proposal submission, and always **requires a very strong scientific justification**. A Spitzer observer may assume that any specific dates and times requested in an approved proposal will be accommodated, unless otherwise notified by the SSC (however, see note in section 5.5.3 below regarding constraint limitations). In the event of a serious conflict, the SSC will contact the affected observer(s) as soon as possible to resolve the problem. However, even for heavily constrained AORs, the only formal announcement of their place in the Spitzer schedule will be their appearance in the latest online weekly schedule and in the above-mentioned e-mail message. We will do our best to maximize the scientific return via scheduling, but specific requests for ordering of AORs in the schedule cannot generally be accommodated.

5.5.3 Requesting Scheduling Constraints for Science Reasons

The scheduling software is able to handle common types of constraints and logical linkages, such as those needed for periodic monitoring of a target. However, because scheduling constraints restrict the ability of the scheduling system to produce efficient schedules and may render scheduling some observations infeasible, **the usage of any type of timing or relational constraint requires a very strong scientific justification**. The justification must be provided at the time of proposal submission (or modification if the constraint has been subsequently added).

5.5.3.1 Recommendations for the Use of Scheduling Constraints

If a scheduling constraint is required, it is **strongly recommended that the minimum constraint necessary** to preserve the scientific goal of the observation(s) be applied. In general, larger timing windows and/or loose groupings are preferred to non-interruptible sequences of observations. For very long observations, strategies which permit independent scheduling of the component AORs are much preferred; this not only enhances scheduling efficiency, but also makes the observation as a whole much more robust against the failure of component AORs. As a general rule of thumb, groups of constrained AORs that occupy more than about half of the time period during which they can be scheduled will have a strong negative impact on efficiency and may not be feasible to schedule. In other words, any constraint that involves a time range should be twice as long as the AORs you want scheduled within that time range. For example, 40 hours of observations that must be done within a 40-hour period cannot be scheduled, whereas the same 40 hours of AORs may be quite feasible if they can be scheduled anytime within the same one-week period. The ability to be scheduled is further enhanced by making the observations shorter; 40 hours of 2-hour AORs constrained to a 50-hour period may be barely feasible, whereas 40 hours of 6-hour AORs cannot be scheduled in a time period shorter than 2 weeks.

When designing scheduling constraints, the observer must take into account the natural 12 to 24-hour time scale for PAOs and that Instrument Campaigns will normally be ~5 to ~21 days long. In addition, PCS calibration and instrument and spacecraft maintenance activities can take 15% or more of the available time during a given PAO, but usually takes just a few percent of the available time. Some PAOs, e.g., at the beginning and end of a campaign, may contain a higher fraction of calibration and maintenance activities.

5.5.3.2 Specification of Scheduling Constraints

Scheduling constraints should be submitted at the time the AORs are submitted, using the constraint-editing capabilities of Spot. Please see the Spot User's Guide for details on how to specify constraints, and the Spitzer Observation Planning

Cookbook for an example of creating constrained observations. Note that constraints are requests that the SSC will work to accommodate. In the event of a serious conflict, the SSC will contact the affected observer(s) as soon as possible to resolve the problem. However, even for heavily constrained AORs, the only formal announcement of their place in the Spitzer schedule will be their appearance in the latest online weekly schedule, as described above. **Any scheduling constraint imposed on an AOR will increase the difficulty of scheduling the AOR. Highly constrained observations are at significant risk of non-execution in Cycle-5.**

The following types of scheduling constraints are supported for both inertial and moving targets:

Timing Constraints
> Timing constraints consist of defining a window or series of windows for the start time of an AOR. Absolute-time observations that will be executed at a specific time, or no more than 3 seconds later, can only be supported for moving targets. To specify an absolute time observation in a *moving target* AOR, set the open and close times for the timing window to be identical. Spitzer's scheduling architecture generally operates on relative time, so for inertial targets, the (inertial target) AORs will simply run in order. Timing constraints for inertial target AORs should be macroscopic (days, weeks, months), not microscopic (seconds, minutes, hours).

Relational Constraints
> Relational constraints are ordering or grouping constraints that are applied to a group of AORs. There are four basic types of relational constraint supported by Spitzer:

Chain = Ordered Non-interruptible Group
> A chain can be thought of as a list of AORs that must be executed consecutively in the order specified and without any other kind of activity intervening. Note that the total time for the entire ordered non-interruptible group cannot exceed the maximum time for an individual AOR; for longer observing sequences, an interruptible group must be used. This type of constraint might be used for an on/off source pair of observations.

Sequence = Ordered Interruptible Group
> A sequence constraint is similar to, but less stringent than, a chain constraint. The AORs will be executed in the order specified and a duration in which they should be completed is specified. The sequence constraint should only be used when the science requires sequential ordering of the AORs. For AORs in which the order of observation is not important, a "group within" constraint (see below) should be used, instead of a sequence constraint.

Group-within
> A group-within constraint specifies that a group of AORs will be executed within a specific length of time, but with no particular starting date/time constraint. Once the first AOR has been executed, the rest of the AORs in

the group will begin within the specified time interval. They may be executed in any order within the time interval. This is similar to a sequence constraint, but the observations may be executed in any order.

Follow-on

This constraint forces a follow-on AOR to be scheduled within a given time after a precursor AOR. It can be thought of as a statement that Follow-On-AOR must be scheduled within Time-Window after the end of Precursor-AOR. The follow-on constraint can be used to prevent early execution of an observation when the success or content of the follow-on is dependent upon the successful execution of a precursor observation. This constraint could also be used for periodic observation of a target where the interval between observations is relatively short (hours to a small number of days). One AOR may serve as the precursor to more than one follow-on, but a follow-on may have only one precursor (e.g. one follow-on constraint can tie together only two AORs).

Shadow

The shadow constraint is a special case of the follow-on constraint, and is used to obtain background measurements for moving targets. The primary AOR is executed as specified. The shadow AOR will be executed to repeat the track of the primary observations. The selected AOR parameters must be identical in the two AORs. The shadow may be executed before or after the primary AOR. Note that the shadow does not re-observe the target at a later date, but rather the background of the primary observation. (As with all constraints, shadow observations must be strongly scientifically justified in the observing proposal.)

Timing and Relational Constraints can be combined. For example, a series of AORs used to obtain spectra of a comet over a long track which needed to be broken up into segments due to curvature, might be constrained as a chain with an associated timing constraint related to the acceptable range of solar elongations.

Spot software checks for viability of relational constraints. Spot will warn the user if a timing constraint is set such that the target is not visible for some portion of time within the timing window, but Spot will still allow the constraint to be set. Observers should make sure that the timing windows they introduce overlap sufficiently with the visibility of the target(s), in order for the complete AOR, or set of relationally constrained AORs, to be completed. *Timing windows that barely overlap a visibility window for a given AOR will make scheduling of that AOR very difficult or perhaps impossible* (e.g., through small visibility window changes caused by a more accurate Spitzer orbit calculation).

5.6 Solar System Objects

Spitzer supports observations of SSOs, tracking in linear segments at rates up to $1''$/s as described in Chapter 3. All instruments and all observing modes can be used while tracking. During IOC/SV, IRS peak-up for IRS spectroscopy was

successful for all moving targets attempted, which included objects with rates between <1.0″/hr and just over 200″/hr, and fluxes between 40 mJy and 1 Jy at 15 µm. Peak-up on both point source and extended targets are supported; during IOC/SV, both moving point sources and extended sources were acquired. However, note that peak-ups are not restricted to be performed only on the moving target being tracked, or a point source that is co-moving with it; instead, observers can peak-up on an inertial target and then offset to a moving target. See the Spot User's Guide for details. Use of PCRS for peak-up on moving targets is an unsupported mode. Updated information on Solar System specific observing issues can be found on the web in the "Solar System Observing with Spitzer" document in the Solar System subsection of the Proposal Kit section on the SSC website.

5.6.1 Tracking Performance

Spitzer's SSO tracking capability is similar to its scanning capability and the performance is described in section 3.4.4.

One significant consequence of this for the SSO observer involves sources whose tracks, on an equatorial map, have significant curvature during an AOR. Such AORs may need to be broken up into a series of short (linear) AORs. The spacecraft does not carry any target ephemerides on board, so the track is defined at the time of scheduling and formulated as a vector rate in an equatorial frame. A start time and equatorial J2000 start point and time are also provided for use by the PCS. Once the track command has been issued, the on-board system maintains knowledge of where the telescope should be at what time, and "catches up" with the specified track and maintains it.

Any further updates to SSO tracking performance will be posted in the Solar System subsection of the Proposal Kit area on the SSC website.

5.6.2 Ephemeris Management

Spitzer uses a database of ephemerides for known SSOs derived from the Horizons database maintained by the Solar System Dynamics group at the Jet Propulsion Laboratory. For proposal planning purposes, Spot can retrieve the ephemerides for a specified target, by resolving the NAIF ID. These ephemerides are used to calculate visibility windows and resource estimate calculations for SSOs through Spot.

If the object is newly discovered, is not released through the Minor Planet Center (MPC), and does not yet have a NAIF ID, the observer will still be able to send the orbital elements to the Spitzer Helpdesk (help@spitzer.caltech.edu) and request generation of an ephemeris for planning purposes. Requests received at the SSC for non-standard ephemerides during the last week of the Proposal Call will be supported on a best effort basis.

Note that observations which require ephemeris updates <2 weeks prior to execution significantly perturb the scheduling process and must be identified in the proposed AORs at the time of submission via the "special" button in the Spot AOT window. Late ephemeris updates are considered equivalent to medium- or high-impact ToOs, depending upon the amount of lead time given (see Spitzer Observing Rule #5), and will incur a special observing overhead (for more details see http://ssc.spitzer.caltech.edu/documents/rules/overheads.html).

5.6.3 Target Specification and Scheduling Considerations

SSO observations will normally be performed as tracked observations with both the position and the track rate/direction computed from the orbital elements derived as discussed in the previous section.

5.6.3.1 Shadow Observations

The infrared flux from background sources, and particularly small-scale structure in that background, frequently limits the sensitivity of Spitzer, particularly in the wavelength range 24–160 µm. To assist in background subtraction, Spitzer Solar System observers can specify "shadow" background observations for all instruments and observing modes. A shadow observation allows the track across the sky taken during observation of a moving target to be replayed or pre-played when the target is not there. Shadow observations allow the Solar System observer to remove background small-scale structure, thereby improving moving target sensitivity. To reduce potential errors due to time-dependent changes in the zodiacal light, instrument characteristics, and calibration, a shadow observation is generally most effective when taken as close in time to the primary observation as is scientifically possible. In addition, the shorter the interval of time between primary and shadow, the larger the visibility window for the primary-shadow pair, and the higher the likelihood that the observation can be scheduled. (Moreover, specifically for IRS long-wavelength observations, this also minimizes the rogue pixel effect; see the IRS chapter for more information.) The Spot visualization tools and instrument sensitivity information can be used to determine whether a particular Solar System observation would benefit from a shadow observation.

As with all observing constraints, the need for a shadow constraint must be scientifically justified. Execution of the shadow observation *after* observation of the science target is the default, although there may be situations where this is not the optimal execution order. However, shadows specified to execute *before* the primary observation are much more difficult to schedule than shadows executed after the primary, and although allowed, require even more compelling scientific justification in the proposal than shadows executed after the primary.

5.6.3.2 Timing Constraints

SSO observations can be scheduled with timing constraints, as long as these constraints are scientifically justified. Observers are encouraged to specify observing windows that are as generous as possible, within the scientific constraints, to both enhance the efficiency of our scheduling system and to

increase the probability that the observation can be scheduled. However, when precise timing is scientifically required, an absolute time can be specified for the SSO observation, and the flight computer will start the track and initiate the data taking at that time, or no more than 3 seconds later. Note that there are other required activities at the beginning of the moving target AOR which could vary in duration from target to target, such as slewing to the target from a previous target, or settling and waiting at the acquisition point before starting the track. However, these variable duration activities are executed *before* the absolute timing constraint is applied, to specify the beginning of the track.

The scheduling constraints that are supported for both inertial targets and SSOs are described in section 5.5.3.

5.6.4 Science Instrument Issues

There are no restrictions on the usage of any of the AOTs for tracked SSO observations.

The saturation of the science instruments is discussed in each instrument-specific chapter, and the scattered light characteristics of the telescope are discussed in Chapter 4, in addition to the instrument-specific chapters.

Please note that since SSOs tend to be near the ecliptic plane, the focal plane orientation with which they can be observed is highly restricted (see section 3.2.5).

5.6.5 Overheads

The time duration of each AOR an observer submits with a program includes a fixed slew overhead. For both inertial and moving targets, this is 215 sec in duration. For Cycle-5, an additional 300 sec overhead is now charged for moving targets. This extra overhead accounts for the target acquisition and tracking requirements for moving targets (see "Solar System Observing with Spitzer" in the Proposal Kit for details).

5.6.6 Best Observing Practices for Solar System Objects

The following "best practices" for Solar System objects are provided *in addition to* the discussions of best practices in the individual instrument chapters. There is also a moving target example in the Observation Planning Cookbook, found in the Proposal Kit on the SSC website.

5.6.6.1 Choosing your targets

Whereas fixed targets can have absolute time errors on the order of an hour (usually executing earlier than expected), moving targets can be executed with an absolute time error of only -0, +3 seconds. Please be aware of this if you are planning to observe a Solar System target as an absolute time-fixed target observation. Generally, it is better to request moving target observations (i.e., with tracking) for Solar System observations.

Remember to check for availability of positional error information for the ephemerides for your targets (instructions in the document entitled "Horizons tips for Solar System Observers," found on the SSC website). Some objects have extremely poor positional information (e.g., several 100s of arcseconds) and would have a very low probability of being acquired by Spitzer. There are 3 known ways of dealing with targets in this situation:

1. Select another target of comparable scientific interest, but with better ephemeris accuracy.

2. Use ground-based support facilities to acquire astrometric observations of your proposed target. Please send these data in MPC ingest format to the Helpdesk (help@spitzer.caltech.edu) and we will ensure that they are included in Horizons.

3. Request that SSC ask the JPL Solar System Dynamics team to provide observational support to improve the object's ephemeris. Note that this capability is limited.

Remember to check for availability of ephemerides in the Spot software as soon as possible. For small target lists, you can enter targets in Spot and check for visibility. For larger target lists, it may be more efficient to check the ASCII list of NAIF IDs that have Spot-readable ephemerides. This list is available on the Proposal Kit section of the SSC website. If your target does not appear in this list, even if you are not sure that you will include a given object in your final proposal, please request ephemerides from the SSC as soon as possible. Do not leave this to the last minute! In the last week before the proposal deadline, we will be providing ephemerides to proposers on a best-effort basis only.

5.6.6.2 Late Ephemeris Updates

Late ephemeris updates are most useful if your object has a high probability of behaving unpredictably *during* the 5 weeks prior to the observation, as might be the case for an active comet near perihelion. Another good justification for a late ephemeris update is a target that, just prior to observation by Spitzer, has relatively few observations and a correspondingly large positional error, as might occur for a newly discovered object that has been requested as a "no-impact" (more than 5 weeks warning of desired observation date) ToO observation. However, if your target is known to have poor positional accuracy at the time you submit your proposal to a Spitzer GO call, it is NOT necessary to request a late ephemeris update. In this case, there is likely to be several months of time in which to improve the ephemeris, but you should be proactive about acquiring and transmitting the additional observations required to reduce the ephemeris error.

5.6.6.3 Timing Constraints

Make sure that any observer-imposed timing constraints are still within the Spitzer visibility period for your target (see section 3.2 above). Even though a

given timing window may be optimal for the science goals, the object may not be visible to Spitzer in that timing window.

Also be aware that not all instruments are available all the time. The Baseline Instrument Campaign (BIC) is published regularly on the SSC website; see section 3.3.1 above. After proposal selection, we may need to iterate with you to select a timing constraint that is good for your science, but that also follows the schedule for instrument availability outlined in the BIC.

5.6.6.4 Visualization

It is EXTREMELY important to visualize your proposed AORs. This is especially true for moving target observations.

Spot will allow you to overlay the orbital track of your object on the infrared background, and will also overlay your AOR in the sequence in which it would be observed on a date that you enter. You can use this feature to check for structure in the infrared background around your object. For IRS observations with peak-up, it can also be used to check along the orbital track over which your object is visible for fixed targets that are brighter than your science target that might appear in the peak-up array and compromise the peak-up process.

If you have a time-constrained observation, you can use the "Animate with Trail" feature in the visualization tool to see the time sequence and positions at which individual apertures will be observed.   For cometary, satellite, or ring observations in particular, check this time sequence for potential unplanned slews across the cometary nucleus, major planet, or some other bright object.

5.7 Targets of Opportunity

ToOs (see Observing Rule #5) are transient events whose timing is unpredictable. Predictable phenomena whose precise timing is not known a priori (e.g., novae, newly discovered comets, gamma-ray bursts) may be requested in a General Observer (GO) proposal. Observations of completely unanticipated phenomena can be requested through DDT procedures (see SSC website for details). A GO proposal must include an AOR for each predictable ToO observation, filled in with as much information as possible.

An example demonstrating proper construction of ToO AORs is included in the Observation Planning Cookbook.

5.7.1 Classification of ToOs

ToOs are classified based solely on their impact on the observatory scheduling process, which depends on the time elapsed between the activation of a ToO observing request and the desired date of execution of the corresponding observation; see Table 5.1 for ToO classification criteria.

Table 5.1: Classification of ToOs

High impact	<1 week (48 hour minimum turnaround)
Medium impact	1–5 weeks
Low impact	>5 weeks

Special overhead burdens will be assessed against high- and medium-impact observations. See this discussion in the Proposal Kit section of the SSC website.

Because of the significant effect that high/medium-impact ToO observations have on efficient telescope scheduling, the combined total of high/medium-impact ToO activations approved in Cycle-5 will not exceed five.

5.7.2 Activation of AORs

A request for ToO activation must be sent to the SSC Director via the Helpdesk (help@spitzer.caltech.edu). If the ToO involves a previously accepted AOR, then the request is to activate this AOR. Modification (e.g., to position) of the AOR is possible, and a request for this should be included with the ToO activation request. If an AOR for the ToO does not yet exist, a request for Director's Discretionary Time (DDT) for the ToO should be made.

Target of Opportunity activation requests will be supported around the clock. Note that activation of a high-impact ToO will take a minimum of 48 hours between receipt of final approved AOR and observation. Further instructions for activating ToOs, including via a request for DDT, are available on the SSC website.

5.8 Generic Targets

Generic Targets (see Observing Rule #6) can be scientifically described, but do not have exact celestial coordinates or brightness estimates at the time of proposal submission. An example of this would be objects discovered in a survey (by Spitzer or any other telescope) for which IRS spectroscopy is sought. Integration time estimates within a factor of 1.5 must be provided in the generic target AOR and a position within 2° of the expected final position. For a moving generic target, the position should be given as 'to be determined.' See the Spitzer Observing Rules for more details. Generic targets should be specified in Spot via the "Special…" button found in the AOT window. **For Cycle-5 generic targets, AORs must be completely specified and ready to schedule by June 1, 2008.**

5.9 Second-Look Observations

In previous cycles, predictable and pre-planned re-visits to objects and/or fields have been an appropriate as part of observing investigations. Because of the shortened cycle, **these 'second-look' observations are not allowed as part of any Cycle-5 proposal.**

6 Infrared Array Camera (IRAC)

The InfraRed Array Camera (IRAC) was built by the NASA Goddard Space Flight Center (GSFC) with management and scientific leadership by the Smithsonian Astrophysical Observatory (SAO) under principal investigator Giovanni Fazio. The information in this manual is based on the design requirements and on the characterization of the flight instrument in pre-flight ground tests and on in-flight performance, including the In-Orbit Checkout (IOC)/Science Verification (SV) period in August-November 2003.

A brief, high-level summary of IRAC for astronomers appeared in the ApJS Spitzer Special Issue, specifically the paper by Fazio et al. (2004, ApJS, 154, 10) entitled "The Infrared Array Camera (IRAC) for the Spitzer Space Telescope" and in the paper by Hora et al. (2004, SPIE, 5487, 244) entitled "In-flight performance and calibration of the Infrared Array Camera (IRAC) for the Spitzer Space Telescope." Copies of these papers are available on the SSC website.

6.1 Instrument Description

6.1.1 Overview

6.1.1.1 Instrument Description

IRAC is a four-channel camera that provides simultaneous $5.2' \times 5.2'$ images at 3.6, 4.5, 5.8, and 8 μm. Two adjacent fields of view are imaged in pairs (3.6 and 5.8 μm; 4.5 and 8.0 μm) using dichroic beamsplitters. All four detector arrays in the camera are 256×256 pixels in size, with a pixel size of $\sim 1.2'' \times 1.2''$. The two short wavelength channels use InSb detector arrays and the two longer wavelength channels use Si:As detectors.

The IRAC instrument was designed to address the four major scientific objectives defining the Spitzer mission. These are (1) to study the early universe, (2) to search for and study brown dwarfs and superplanets, (3) to study ultraluminous galaxies and active galactic nuclei, and (4) to discover and study protoplanetary and planetary debris disks. The utility of IRAC is in no way limited to these objectives, which we only mention to explain the scientific drivers for the instrument design. IRAC is a powerful survey instrument because of its high sensitivity, large field of view, mapping capabilities, and simultaneous four-color imaging.

6.1.1.2 Technical Overview

IRAC consists of the Cryogenic Assembly (CA) installed in the Multiple Instrument Chamber (MIC) in the CTA, and the Warm Electronics Assembly (WEA) mounted in the spacecraft. Harnesses connect the detectors and calibration subsystem in the CA to the WEA. The WEA communicates with the spacecraft over three RS-422 serial lines that allow receiving commands from, and sending acknowledgements and image data to, the spacecraft Command & Data Handling (C&DH) computer.

The IRAC Cryogenic Assembly, depicted in Figure 6.1, consists of the following major subassemblies: the Pickoff Mirrors; the Shutter; the Optics Housings, which hold the doublet lenses, beamsplitters, filters, and cold stops; the Focal Plane Assemblies (FPAs) that include the detector arrays and associated components; the Transmission Calibrator with its Source and Integrating Spheres; and the Housing Structure, consisting of the Main Housing Assembly and the wedge-shaped MIC Adapter Plate.
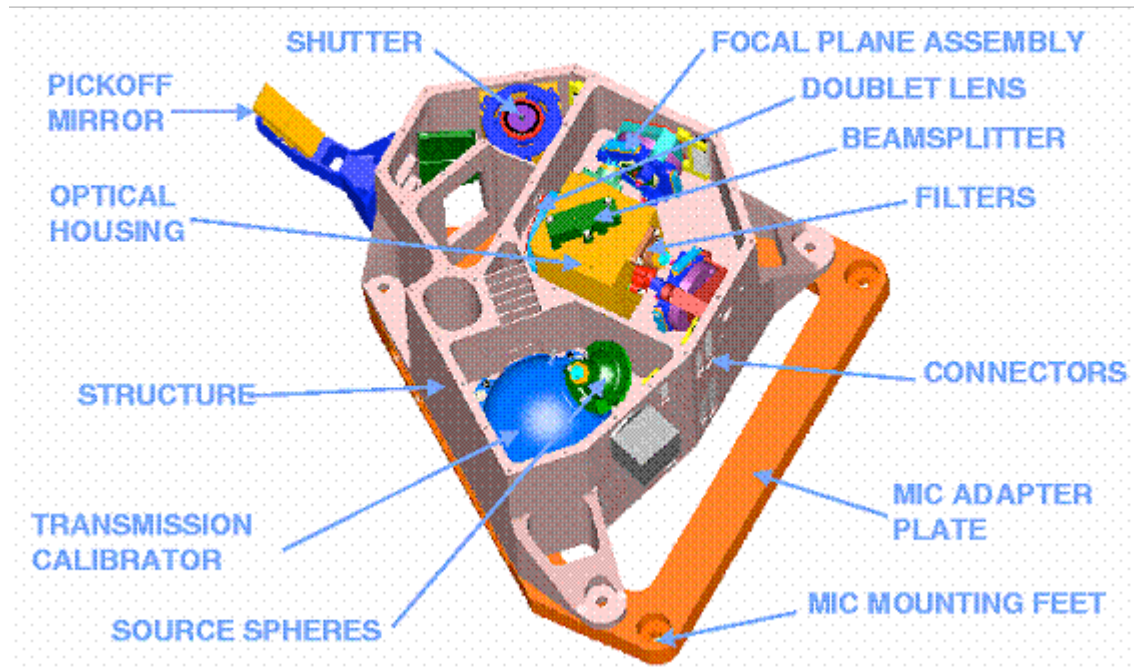
Figure 6.1: IRAC Cryogenic Assembly model, with the top cover removed to show the inner components.

6.1.2 Optics

6.1.2.1 Optical Layout

The IRAC optical layout is shown in Figure 6.2 and Figure 6.3. Light from the telescope is reflected into the IRAC structure by the pickoff mirrors for the two fields of view (FOVs). Each pair of channels has a doublet lens which re-images the Spitzer focal plane onto the detectors. A beamsplitter reflects the short wavelength light to the InSb detectors (Channels 1 and 2) and transmits the longer wavelength light to the Si:As detectors (Channels 3 and 4). Channels 1 and 3 view the same telescope field (within a few pixels), and Channels 2 and 4 view a different field simultaneously (see Figure 4.5– the Spitzer FOV). However, the orientation of the sky on the InSb arrays are flipped with respect to the Si:As arrays by the beamsplitter. The edges of the two IRAC fields of view are

separated by approximately 1.5′, with no overlap on the sky. The IRAC pixel scale is nearly the same in all channels (~1.2″ per pixel), providing a 5.2′×5.2′ FOV.
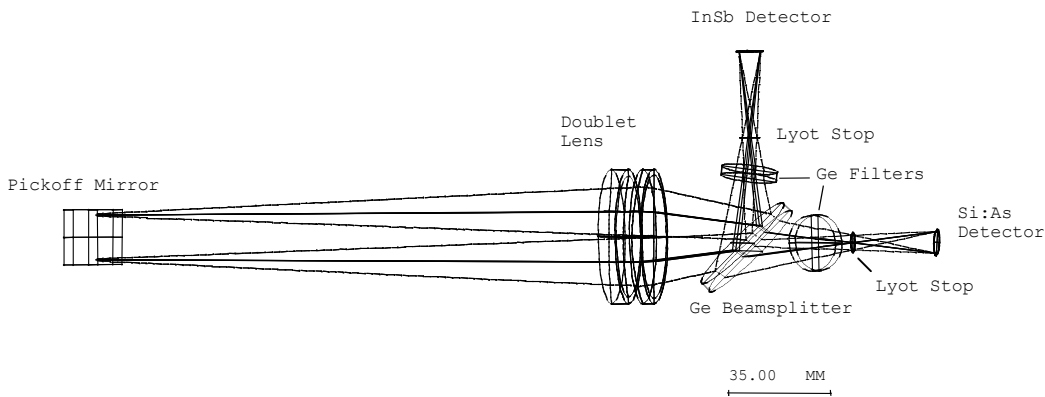
Figure 6.2: IRAC optical layout, top view. The layout is similar for both pairs of channels; the light enters the doublet and the long wavelength passes through the beamsplitter to the Si:As detector (Channels 3 and 4) and the short wavelength light is reflected to the InSb detector (Channels 1 and 2).
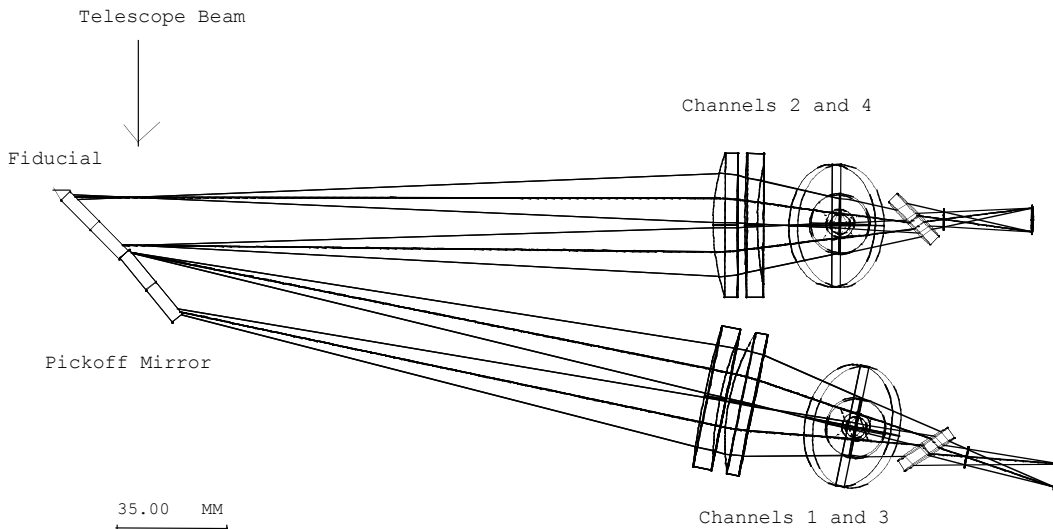
Figure 6.3: IRAC optics, side view. The Si:As detectors are shown at the far right of the figure, the InSb arrays are behind the beamsplitters.

6.1.2.2 Image Quality

The IRAC optics specifications limit the wavefront errors to $< \lambda/20$ in each channel. IRAC provides diffraction-limited imaging internally, and 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 effect is

small, with the total lateral chromatic dispersion less than a pixel in the worst case. The sky coordinates of each pixel have been accurately measured in flight using astrometric observations of an open cluster, resulting in distortion coefficients that are in the world coordinate system of each image. 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.

A much larger variation in the flux of sources measured in different parts of the array is due to the tilt of the filters, which leads to a different spectral response in different parts of the field of view. The flat field calibration is done with the zodiacal light, which is relatively red; blue sources have a flux variation of up to 10% from one side of an array to the other (see the IRAC Data Handbook for more details).
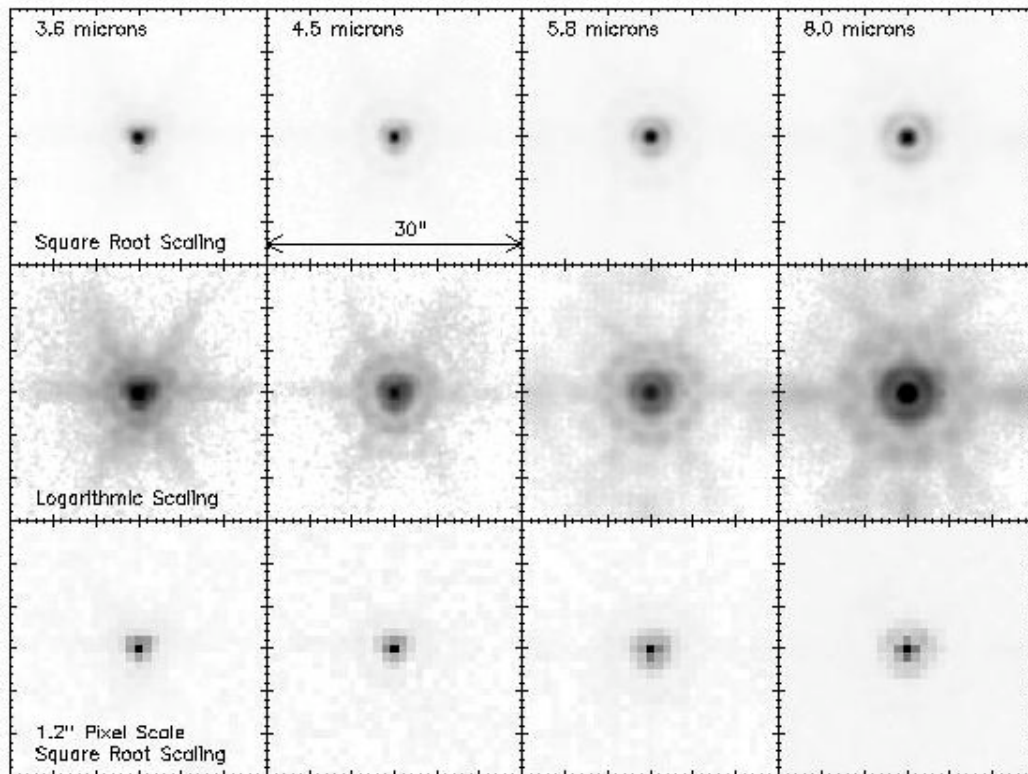
Figure 6.4: The in-flight IRAC point response functions (PRFs) at 3.6, 4.5, 5.8 and 8.0 μm. The PRFs were reconstructed onto a grid of 0.3″ pixels, ¼ the size of the IRAC pixel, using the drizzle algorithm. We display the PRF with both a square root and logarithmic scaling, to emphasize the structure in the core and wings of the PRF, respectively. We also show the PRF as it appears at the IRAC pixel scale of 1.2″. The reconstructed images clearly show the first and second Airy rings, with the first Airy ring blending with the core in the 3.6 and 4.5 μm data.

Figure 6.4 shows the IRAC point response functions (PRF) reconstructed from images of a bright star obtained during IOC/SV. (Here we use the language common in the optics field; the point spread function [PSF] is before sampling by

the detector array, and the point response function [PRF] is after sampling by the detector array. More information is available on the SSC website.) The PRFs, which are undersampled at the IRAC pixel scale, were generated by combining 108 individual IRAC images in each band. By offsetting each image by a fraction of a pixel width, fully sampled PRFs can be extracted from the data. The resulting PRFs are the "optical" point spread function projected onto the focal plane by the IRAC and telescope optics, convolved with the response function of a single detector pixel. The images were combined using a drizzle algorithm (Fruchter & Hook, 2002, PASP, 114, 144) to minimize smoothing of the PRF during the reconstruction process. The resulting pixel scale was ¼ the width of an IRAC pixel. Images of a bright star at the native IRAC pixel scale are also displayed for comparison. FITS images of IRAC PRFs are available on the SSC website. The appropriateness of a given PRF is dependent on the observation sampling and the photometric reduction package used (see IRAC Data Handbook).

Table 6.1: IRAC image quality properties

Channel	Noise pixels (mean)	FWHM (mean;")	FWHM of centered PRF (")	Central pixel flux (peak; %)	Pixel size (")	Maximum distortion (pixels relative to square grid)
1	7.0	1.66	1.44	42	1.221	1.3
2	7.2	1.72	1.43	43	1.213	1.6
3	10.8	1.88	1.49	29	1.222	1.4
4	13.4	1.98	1.71	22	1.220	2.2

Table 6.1 shows some properties of the IRAC PRF. These numbers were derived from in-flight measurements of bright stars. The PRF was monitored throughout the early mission while the telescope cooled down, and the secondary mirror was moved in two steps to achieve a near-optimal PRF. The "noise pixels" column in Table 6.1 gives the equivalent number of pixels whose noise contributes to a linear least-squares extraction of the flux of a point source from a 13×13 pixel portion of an unconfused image and assuming the PRF is perfectly known. In more detail, the quantity is derived as follows.

Let the PRF in pixel i be P_i and the intensity of an image in pixel i be I_i. If a point source with flux F is present in the image, then $I_i = FP_i$. If we do a least-squares

$$\chi^2 = \sum \frac{(I_i - FP_i)^2}{\sigma_i^2}$$

fit to determine F, then we minimize   where σ_i is the measurement uncertainty in pixel i. We will assume here that σ_i is independent of pixel and set $\sigma_i = \sigma$. Now we take the derivative of χ^2 with respect to the

source flux and set it to zero to find the optimum value. We find

$$0 = \sum (I_i - FP_i)P_i \text{ ; solving for } F, \text{ we find } F = \frac{\sum I_i P_i}{\sum P_i^2}.$$

Now we derive the uncertainty in the flux. Using the well-known theorem for propagation of errors $\sigma_F^2 = \sum (dF/dI_i)^2 \sigma^2$, and applying it to the result above,

we find that $\sigma_F^2 = \sum \left(\frac{P_i}{\sum P_i^2} \right)^2 \sigma^2 = \frac{\sum P_i^2 \sigma^2}{\left(\sum P_i^2 \right)^2} = \frac{\sigma^2}{\sum P_i^2}$, or equivalently,

$\sigma_F = \sigma \sqrt{N}$ where $N = \frac{1}{\Sigma P_i^2}$, which is the definition of noise pixels. The

number of noise pixels is used for the sensitivity predictions in the following sections of this chapter.

There are two columns for the full width at half-maximum (FWHM) of the PRF. The mean FWHM is from observations of a star at 25 different locations on the array. The FWHM for "centered PRF" is for cases where the star was most closely centered in a pixel. The fifth column in Table 6.1 is the fraction of the flux in the central pixel for a source that is well centered in a pixel. It was determined from the images of the focus star (after the telescope was focused) that were the most symmetric and concentrated. These values for the flux in the central pixel were used in the saturation predictions in the following sections of this document. The flux in the central pixel for a random observation will be lower, depending on the centering of the source with respect to a pixel.

The sixth and seventh columns in Table 6.1 were determined from astrometric measurements of star fields.

6.1.2.2.1   Distortion

Due to the off-axis placement of IRAC in the Spitzer focal plane, there is a small amount of distortion over the IRAC FOV. The maximum distortion in each IRAC band is <2.2 pixels (compared to a perfectly regular grid) over the full FOV. Figure 6.5 shows the distortion across all four IRAC channels, as determined from data taken during IOC/SV.

6.1.2.2.2   Location of sources on the arrays

For an observation specifying a target with no offsets or dithers, the source will be positioned at the pointing center of each field of view, modulo the blind pointing uncertainty of Spitzer (~0.5″ rms). In BCD coordinates, the source will appear as listed in Table 6.2. The x and y are the row and column positions, respectively, in the BCD, counting from the lower left-hand corner (1, 1), with integer values specifying the center of a pixel. Offsets in the positive column direction will move the source in the negative y-direction on the BCD. Offsets in the positive row direction will move the source in the positive x-direction on the BCD.
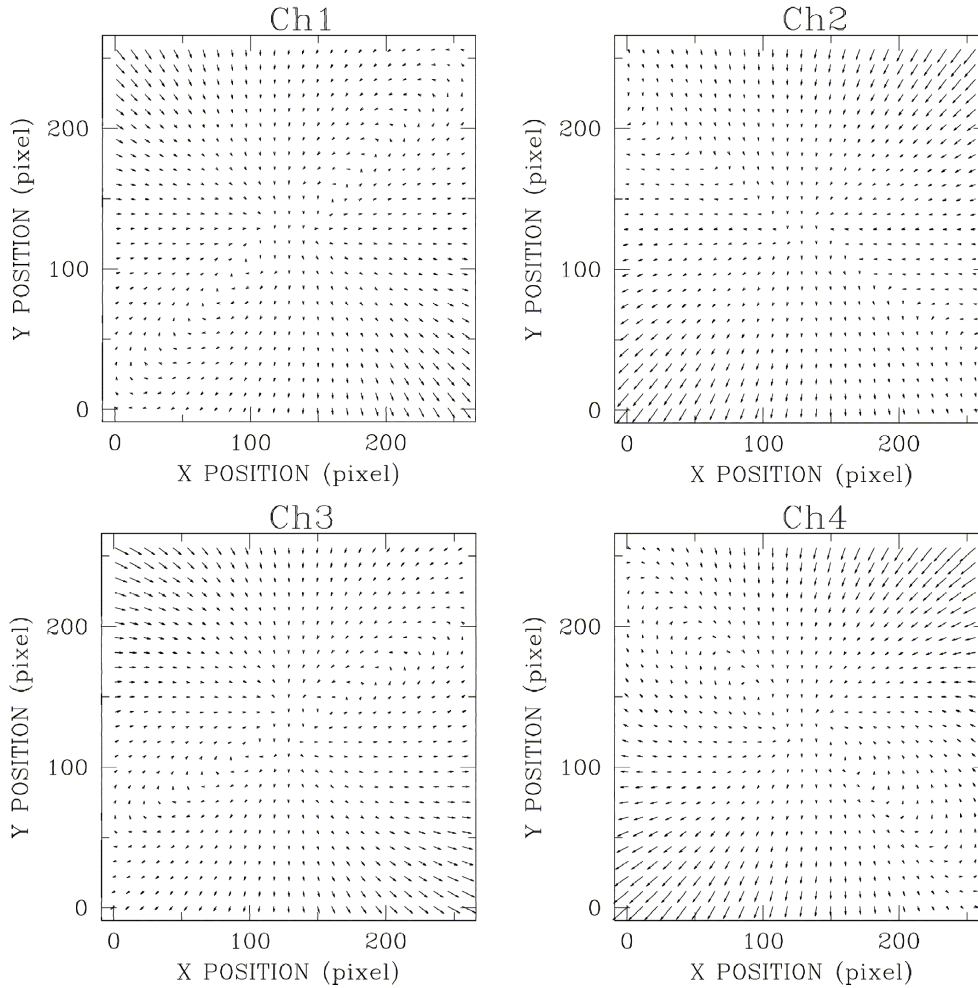
Figure 6.5: Optical image distortion in IRAC channels. The panels show the image distortions as calculated from a quadratic polynomial model that has been fit to in-flight data. The magnitude of the distortion and the direction to which objects have moved from their ideal tangential plane projected positions is shown with arrows. The length of the arrows has been increased by a factor of ten for clarity. The maximum positional deviations across the arrays for this quadratic distortion model are less than (1.3, 1.6, 1.4, 2.2) pixels for channels (1, 2, 3, 4). The derivation of the pixel scales that are listed in Table 6.1 fully accounted for the quadratic distortion effects shown here.

Table 6.2: Source location on the arrays

Array	x (pixel)	y (pixel)
3.6/5.8 μm field of view		
3.6 μm	131.0	128.0
5.8 μm	124.6	128.8
4.5/8.0 μm field of view		
4.5 μm	128.0	129.0
8.0 μm	127.1	127.5

The orientation of the BCDs on the sky is shown in Figure 6.6. The figure also displays the positioning of the subarray fields of view with respect to the full array. The 3.6 μm subarray subtends the location (9, 217) to (40, 248) of the 3.6 μm full array BCD. The 5.8 μm subarray subtends the location (9, 9) to (40, 40) of the 5.8 μm full array. The 4.5 μm subarray orientation with respect to the full array is the same as the 3.6 μm array. Likewise, the 8 μm subarray has the same orientation as the 5.8 μm array.
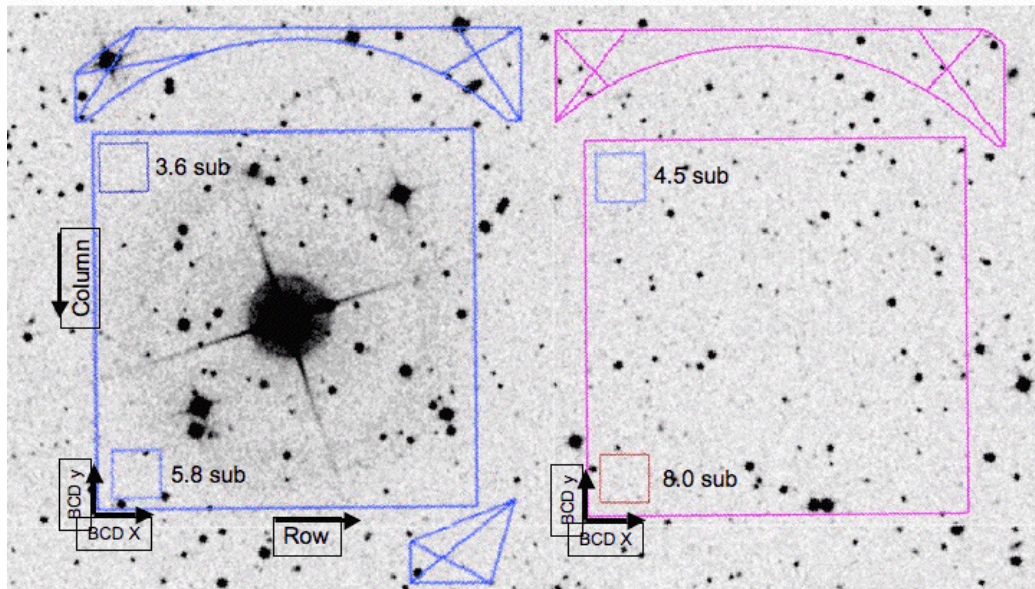
Figure 6.6: The orientation of the BCDs on the sky.

6.1.2.2.3   Scattered and stray light

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 flat fields when they are derived from observations of the sky.

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. The zones for channels 1 and 2 are shown in the Spot overlays (Figure 6.7), as these produce the largest amount of scattered light. 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 (Figure 6.8). They arise due to scattering of light from the edges of the holes in the covers of the focal plane arrays. The A and B zones join up to form a continuous region of scattering which is widest at the upper left and right corners

of the array, and narrows to a waist in the center. Zone 1C produces diffuse scattered light across the array, but the light is only noticeable if Zone 1C is illuminated by very bright stars. See the SSC website for more details. The zones for channels 3 and 4 are a narrow strip about 3 pixels wide, 16 pixels outside of the array and surrounding it. They are produced by a different mechanism involving scattering with the detectors themselves.

Scattered light is well-rejected by the post-BCD pipeline if dithering and mapping offsets are large enough to ensure stars are moved out of the scattering zones, and the redundancy is high enough to allow for effective outlier rejection. The medium and large dither pattern scales with 4 or more dithers per position are adequate for this. Most cases of scattered light are identified and flagged as part of the pipeline processing.
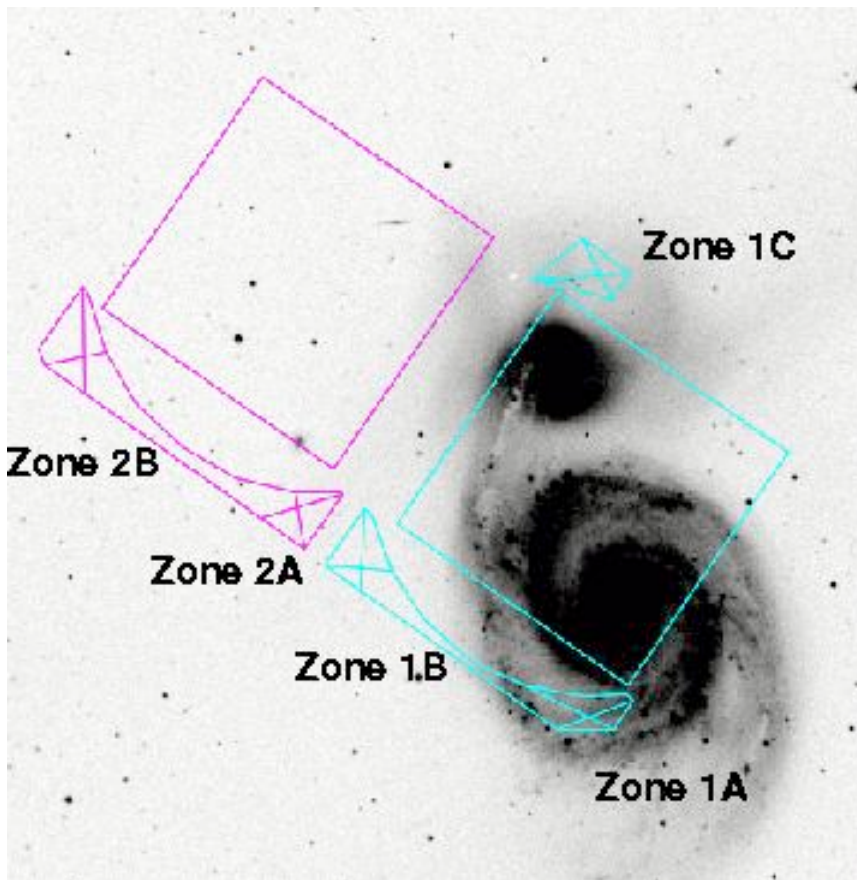
Figure 6.7: Locations of IRAC stray light avoidance zones as seen in Spot visualizations. See text and SSC website for more detail.

The effects of the scattered light on the flats are illustrated in additional figures found on the SSC website. The SSC has implemented a correction in the pipeline to remove diffuse stray light from the BCDs, flats, and skydarks. This correction scales a template of the diffuse scattered light pattern by the expected zodiacal background level and subtracts it from the data. The behavior of the scattered light is different in channels 1 and 2 compared with channels 3 and 4. In channels

1 and 2, the pattern is in the form of a "butterfly wing" at the bottom of the array in raw data coordinates (i.e., at the top of the array in BCD coordinates). The scattered light avoidance zones 1A, 1B, 2A, and 2B are larger than the regions the scattered light falls into. This results in the "butterfly wings" having amplitudes ~5% of the background intensity, even though only about 2% of the light from a given point in the stray light zones is scattered into them. In channels 3 and 4, the pattern consists of narrow strips about 30 pixels from the edges of the array (the "lattice" or "tic-tac-toe" pattern). There is no strip in the flat field on the negative x side of the channel 3 frame where the array is vignetted by the pickoff mirror, and the strip at the top of the channel 3 array is brighter and has a relatively bright secondary reflection about 50 pixels from the top of the array (see Figure 6.13). In channels 3 and 4, the pattern is only ~1% of the background.

Additional information and figures on this subject are available on the SSC website and in the IRAC Data Handbook.
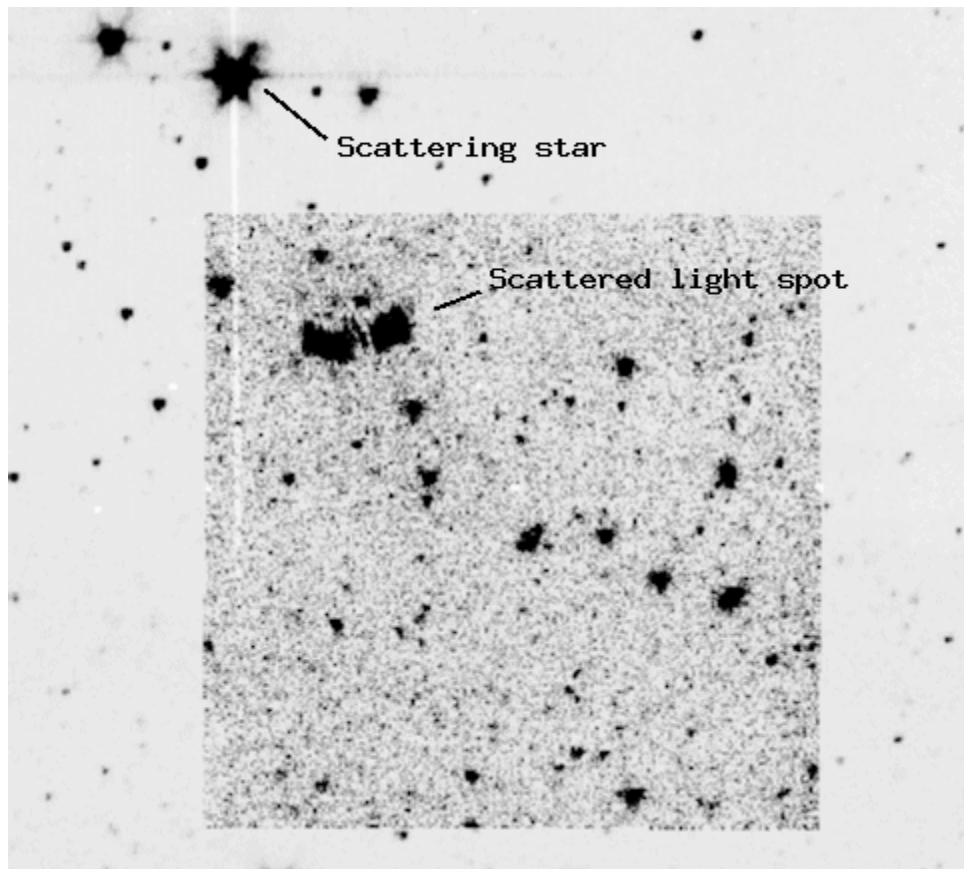
Figure 6.8: Stray light pattern in BCD data when a bright source falls in the stray-light avoidance zone 2A in channel 2. See text and SSC website for more detail.

6.1.2.2.4   Ghost images

Ghost images are visible near very bright sources in channels 1 and 2. These ghost images are caused by internal reflections within the tilted filters. Because of

the increase in the optical path length, ghost images are not in focus. The separation between the main image and its ghost is roughly proportional to the distance of the main image from the Spitzer optical axis in both Y and Z directions (see Figure 4.5), i.e., $(\Delta Y, \Delta Z) = (A_y y + B_y, A_z z + B_z)$ where *(y,z)* are normalized coordinates in which the FPAs span the range [0,1] with the axes increasing away from the Spitzer optical axis. The coefficients for channel 1 are A_y=0.04351, B_y=0.00288, A_z=0.04761, and B_z=0.00211. The coefficients for channel 2 are A_y=0.04956, B_y=0.00105, A_z=0.04964, and B_z=0.000387. The peak intensity of the ghost is roughly 0.05% of the (unsaturated) peak intensity of the star.

Filter ghosts also appear for very bright (saturated) sources in channels 3 and 4. These ghosts are cross-shaped ("+"). They are offset from the primary image by approximately (+36 pix, +2 pix) and (-36 pix, +2 pix) in the Spitzer (Y,Z) directions for channels 3 and 4, respectively. The Z-offset varies slightly with position on the array. The peak intensities of these ghosts are < 0.5% of the (unsaturated) peak intensity of the source; they are in the process of being characterized. Examples of ghost images are available on the SSC website. The SSC pipeline flags filter ghosts for channels 1 and 2.

6.1.2.2.5   Polarization

The tilted elements within IRAC produce a small amount of instrumental polarization that can lead to small photometric errors when viewing polarized sources. The filter transmission is shown as a function of wavelength in Figure 6.9; the expected polarizations are 0.794%, 0.282%, 3.87%, and 2.79% for channels 1–4, respectively.

Table 6.3: IRAC Channel characteristics

Chan	Effective λ (μm)	Bandwidth (μm)	Average transmission (η_I)	Minimum in-band trans-mission	Peak trans-mission
1	3.550	0.750 (21%)	0.676	0.563	0.748
2	4.493	1.015 (23%)	0.731	0.540	0.859
3	5.731	1.425 (25%)	0.589	0.522	0.653
4	7.872	2.905 (36%)	0.556	0.450	0.637

6.1.2.3 Chromatic Elements

The IRAC system throughput and optical performance is governed by a combination of the system components, including the lenses, beamsplitters, filters, mirrors, and detectors. The system parameters are summarized in Table 6.3 and Figure 6.9. The system response is based on measurements of the final in-flight system, including the beamsplitter, filter, ZnS & ZnSe coating transmissions, mirror reflectance, BaF_2 and MgF_2 coating transmissions, and detector quantum efficiency.

The system out-of-band transmission (more than 2.5 times the central wavelength of the band) is less than 10^{-8} relative to the transmission at the center of the band. Tests during IOC/SV showed that the out-of-band leaks are less than the astronomical background at all locations for sources of any temperature detectable in the IRAC bands.

6.1.3 Detectors

6.1.3.1 Physical Characteristics

The IRAC detector arrays were developed by the Raytheon/Santa Barbara Research Center (SBRC) in Goleta, CA, under contract to SAO (Hoffman et al. 1998, Proc. SPIE, 3354, 24; Estrada et al. 1998, Proc. SPIE, 3354, 99). Channels 1 and 2 use InSb arrays operating at ~15 K, and channels 3 and 4 use Si:As detectors operating at ~6 K. Both array types use the CRC744 CMOS readout circuit, and the same physical pixel size of 30 μm. The arrays are anti-reflection coated with SiO (Channels 1, 2, and 3) and ZnS (Channel 4). The power dissipation for each array is <1 mW. Table 6.4 gives some of the detector properties for IRAC channels 1–4. The "operability" is the percentage of the pixels in an array that are within usable specifications.
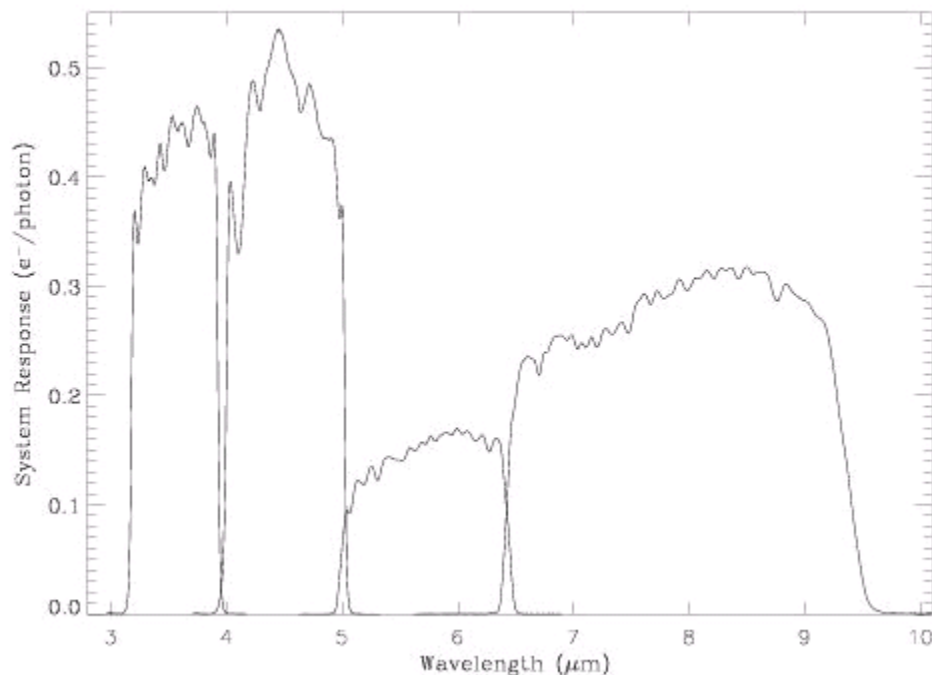
Figure 6.9: IRAC total system response, including transmission of the optics and quantum efficiency of the detectors. Digitized versions of these curves are available on the SSC website.

Table 6.4: IRAC detector characteristics

	FPA designation	Read noise for frame with specified frame time (electrons)					Quantum Efficiency (%) •	Well Depth (e-) •	Operability (%) •
		0.1s*	2s	12s	30s	100s			
1	48534/34 (UR)	16.9	11.8	9.4	7.8	8.4	87	145,000	99.97
2	48975/66 (GSFC)	16.8	12.1	9.4	7.5	7.9	86	140,000	99.9
3	30052/41 (ARC)	9.0	9.1	8.8	10.7	13.1	45	170,000	99.99
4	30219/64 (ARC)	8.4	7.1	6.7	6.9	6.8**	70	200,000	99.75

* Per single subframe (1 of the 64 planes in the BCD cube).
** Per 50 s frame.

6.1.3.2 Detector Sensitivity Calibration

6.1.3.2.1 Photometric linearity calibration

Both types of detectors have measurable nonlinearity. The InSb arrays are nearly linear until they reach saturation. The Si:As detectors are somewhat nonlinear over most of their operating range, and above half-well capacity this will contribute noticeably to the total error budget. However, all of the arrays are linearized to better than 1% up to approximately 90% of their full-well capacity (defined in electrons in Table 6.4, with the gain listed in Table 6.6, corresponding typically to 45,000-60,000 DN). The detector linearity has been measured during ground testing and in flight. The laboratory linearity measurements, with the flight instrument, are shown in Figure 6.10. The arrays were illuminated with a constant flux, and successively longer exposures were taken. For a perfectly linear system, the flux would be directly proportional to the exposure time, and the graph would show a straight line. In fact, the arrays were driven past their saturation levels, and the shape of the curve up to 90% of the saturation level was fitted with a polynomial for the linearization module in the pipeline.

6.1.3.2.2 Effects of overexposures

After observing a bright source, a residual image will remain on the array in subsequent frames. In pre-launch laboratory testing, we found that after taking a 30 sec frame on a very bright lamp, the residual image in the next 30 sec frame had an amplitude (relative to full well) of 0.6%, 1.1%, 2.1%, and 3.7% in channels 1, 2, 3, and 4, respectively. The residual image decayed and became indistinguishable within 270 sec. Exposing the arrays to lamps that lightly saturated them led to residual images of lower amplitude: in channels 1 and 2, a lightly saturating source produced no strong residual image, and the amplitudes of the residual images in channels 3 and 4 were only 0.3% of full well. The predicted saturation levels for point sources are given in Table 6.13.
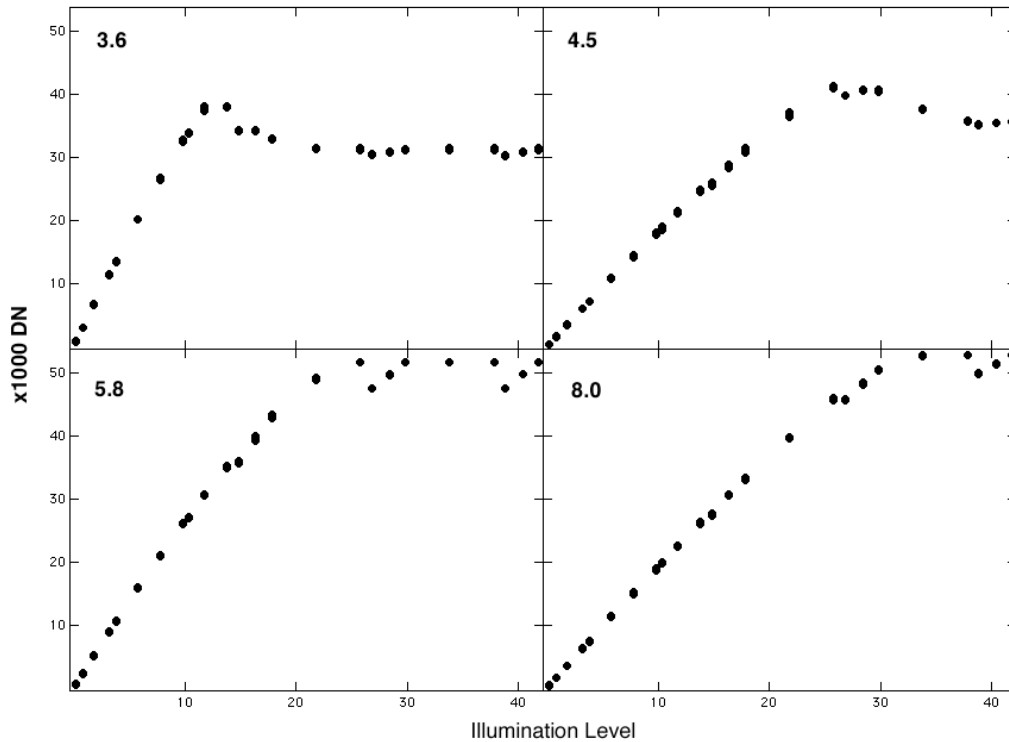
Figure 6.10: Non-linearity curves for the IRAC detectors. The detector responses are fairly linear until saturation, where there is a steep drop-off in responsivity.

In flight, we discovered that bright sources have very long-lasting effects in channels 1 and 4. These "persistent" images are qualitatively different from the short-term "residual" images, which are described in the next subsection and were reasonably well characterized before launch. In channel 1, sources with 2MASS Ks magnitude < 12, exposed for more than 10 minutes, lead to persistent images that last 6 hours. Such exposures can occur when the spacecraft is staring in the Earth-point configuration during downlinks. (The first downlink on the high-gain antenna generated some of the worst persistent images observed to date, because IRAC was pointed in the Galactic plane for an hour.) Observations using the IRAC AOT can also generate persistent images in channel 1. Tests conducted during IOC/SV demonstrated that short exposures (2 sec) generate only very weak persistent images, even when observing sources with K~0. Exposures between 2 sec and 10 min generate intermediate-amplitude persistent images; for an L=12 star, the persistent images in a 100 sec frame are about L=19.5, or 1000 times fainter. We have also found that slews over very bright targets leave persistent images appearing as straight, faint streaks stretching across the array. All these persistent images in channel 1 decay with time, are proportional to the offending source brightness, and have a shape similar to the logarithm of the direct point spread function. Laboratory tests on flight spare arrays have now demonstrated that the persistent images in channel 1 are due to a known defect of the flight-selected array (a broken bias clamp line). To remove persistent images, we thermally anneal the channel 1 and 4 arrays after downlinks. Persistent

images are effectively erased by anneals; all four arrays are annealed at the start of a campaign.

Persistent images in channel 4 can be generated by stars with K_s<6 and infrared sources with comparable 8 μm flux. Unlike the channel 1 persistent images, those in channel 4 do not decay exponentially; instead, they can last for two weeks. Furthermore, the channel 4 persistent images appear at roughly the same raw DN level, independent of frame time, so they have a worse effect on short frames than on long frames. Channel 4 persistent images are generated by normal scientific or calibration observations with the AOT. We have not seen channel 4 persistent images resulting from stares during downlinks. In nominal operations, we are thermally annealing the channel 4 array at the same time as channel 1, so that persistent images cannot build up. All skydark and flatfield observations will take place after anneals, so they should be free of persistent images.

While it is not always possible, steps should be taken to avoid generating persistent images, including keeping bright sources off the arrays and using multiple short exposures instead of longer ones. Also, observations should be planned so as not to be susceptible to persistent images generated by preceding observations. Well-dithered observations will yield mosaics with much-reduced persistent images and are strongly recommended; in straight averages, the amplitude decreases as the number of dithers is increased, and the worst ones are eliminated from robust averages or medians.

6.1.3.2.3   Short-term residual images

All the IRAC channels have residual images of a source after it has been moved off a pixel. When a pixel is illuminated, a small fraction of the photoelectrons is trapped. The traps have characteristic decay rates, and can release a hole or electron that accumulates on the integrating node long after the illumination has ceased. In all arrays, the longest e-folding decay time is about 1000 sec. Even in a long integration that begins shortly after the illumination ends, the residual image never exceeds 1% of the illuminating source. Usually the residual will be orders of magnitude smaller, but the residual image of a star can persist as a false point source for tens of minutes. In channel 1, with a low background, it takes ~ 15 minutes for a persistent image of a bright source to fall below the 3-σ point source detection threshold in 100 sec images. The time to decay below the threshold decreases with longer wavelength, and in channel 4 it is relatively short, because of the high background.

Figure 6.11 shows an example of a sequence of 200 sec frames in channel 2 following a 400 sec integration on a bright lamp (9129 electrons/sec); the pixel wells were strongly saturated. After an hour, the residual image is still present, but it is well below the noise per pixel with a low background in flight, which is about 40 electrons rms in a 200 sec frame. Residual images were measured with a range of integration times, from 2 sec to 200 sec, to sample the fast and slow decays. The residuals are fitted very closely by a model with 5 "species" of traps,

which have distinct exponential decay times (in InSb) of about 0.7, 5, 30, 180, and 1000 sec.
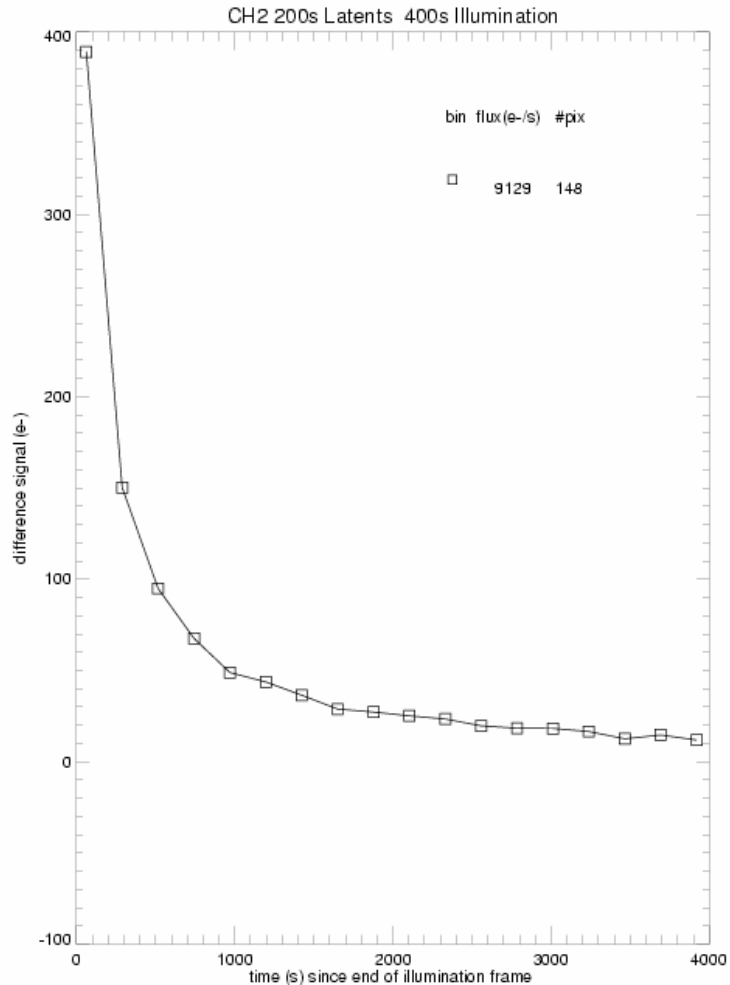
Figure 6.11: Residual image brightness decay as a function of time interval since exposure to a bright source (channel 2).

To illustrate how residual images may appear in an astronomical observation, Figure 6.12 shows a sequence of images of a bright laboratory source in the IRAC channel 1 array. The laboratory source contains optics intended to make the source look like an astronomical target in the far field, but the width of the PSF in the laboratory setup was significantly larger than the PSF that was observed in flight. An observation with relatively small pointing offsets (such as dithers or map grid slews much smaller than the array size) would yield a pattern similar to the one shown here. Each frame shows the bright, primary source together with the residual image of that source from the previous pointings.

6.1.3.2.4   *Pixel-to-pixel sensitivity variations*

The gain varies from pixel to pixel on the array, due to intrinsic gain differences in the pixels themselves or in the output electronics gain, or due to variations in

the light over the FOV transmitted by the telescope and camera optics. Images of the array responsivity are shown for each IRAC detector in Figure 6.13.
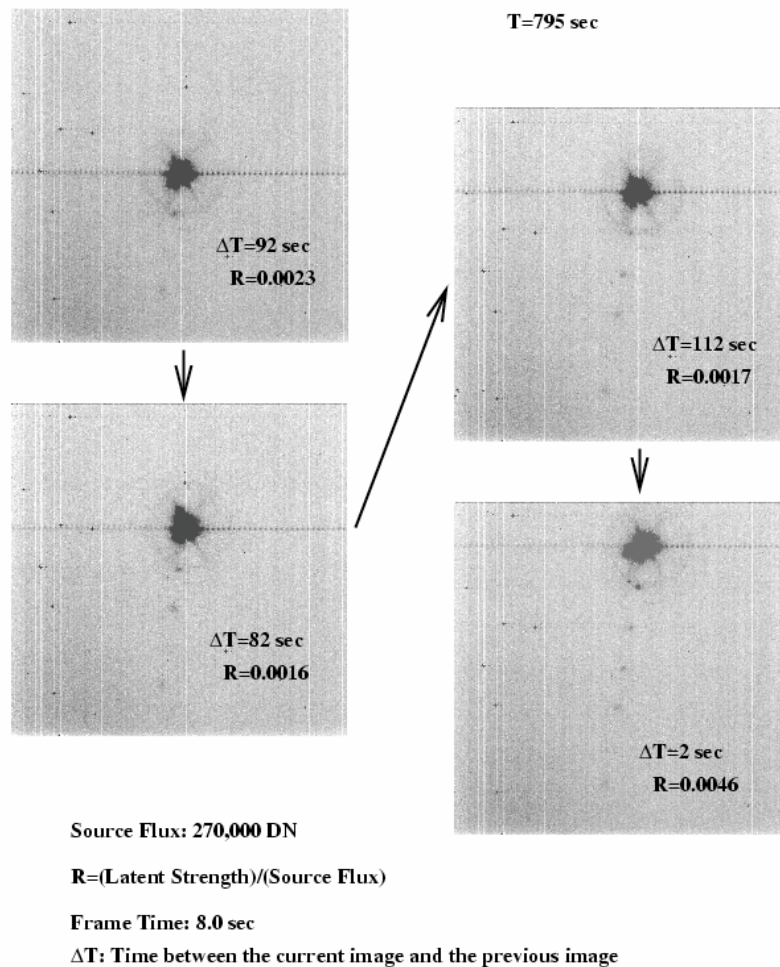
Figure 6.12: Illustration of residual images in lab test data. A bright point source was moved across the array in discrete steps. Each frame shows the bright source (with mux-bleed trails running horizontally) as well as faint residual images from the previous exposures. The interval between frames and the fractional amplitude of the residual are listed for each frame. These are "raw" IRAC images in channel 1.

6.1.3.2.5   Pixel mask

For the InSb detectors, the measured flux of a point source depends on the position of the source centroid with respect to the center of a detector pixel. This effect is due to the efficiency of the InSb pixels not being uniform throughout the entire pixel. The variation in measured flux is ~4% at 3.6 μm and <1% at 4.5 μm comparing a source centered on a pixel to centered the corner of a pixel.

6.1.3.2.6   Intra-pixel sensitivity variations

In each array there are pixels that are totally inoperable, or whose characteristics (QE, noise, dark current) fall outside of the allowed range. This is represented by

the Operability column in Table 6.4. These pixels are evaluated and masked out during reduction in the pipeline using pixel masks. Interested observers can download the quasi-static bad pixel masks from
http://ssc.spitzer.caltech.edu/irac/products/pmask.html.
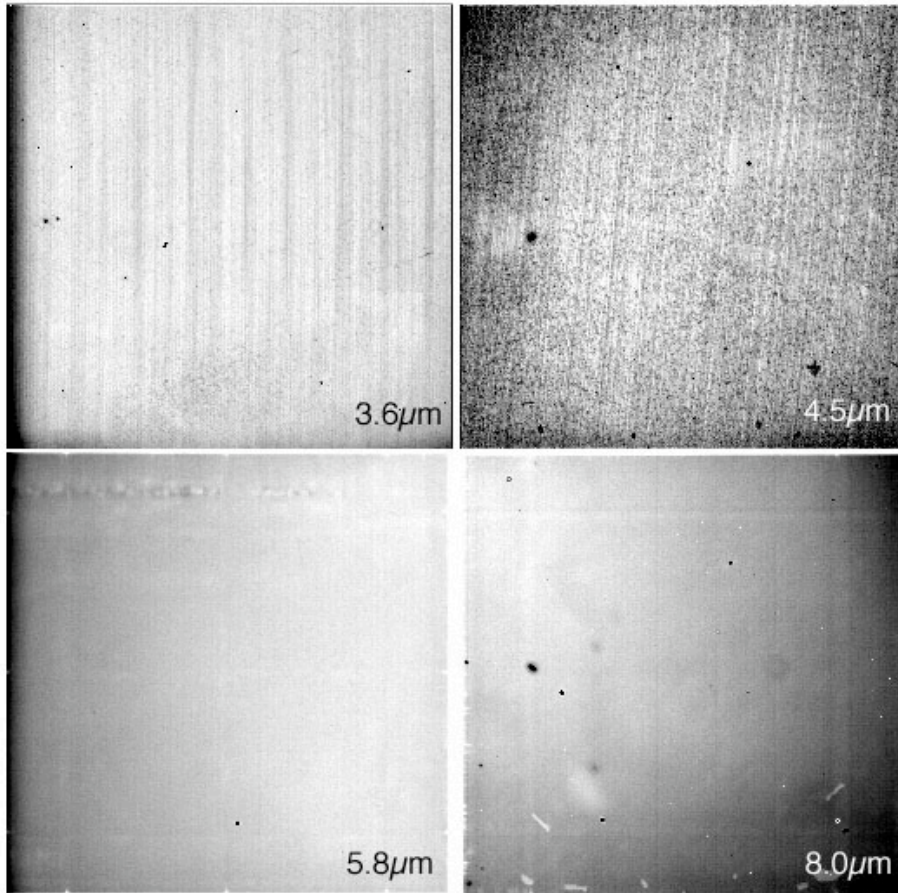
Figure 6.13: IRAC super skyflats for channels 1-4 taken during the first year of nominal operations. The extended dark spots seen in channel 2 and 4 flats are most likely due to particulate impurities on the pick-off mirror.

6.1.3.2.7   Darks

The detector dark currents are generally insignificant compared to the sky background. However, there is a significant offset or bias (which can be positive or negative) in a dark frame, and which therefore must be subtracted from the observations. Figure 6.14 shows dark frames obtained with the flight arrays, using 200 second frames. The average dark current is <0.1, 0.28, 1, and 3.8 e-/s for channels 1, 2, 3, and 4, respectively. Especially for shorter frames, the "dark" images are mostly due to electronic bias differences, rather than true dark current; therefore, the number of electrons in a dark image does not scale linearly with exposure time. Note that without the shutter, isolated dark/bias data (as shown in Figure 6.14) cannot be taken in flight. Instead, lab darks and skydarks are subtracted.

6.1.3.2.8   Flats

Flats are obtained from observations of the sky during normal operations. If the data are taken with the proper dither pattern, it will be possible to relate the total response of each pixel to that of all others. Observations through the entire telescope and instrument are necessary to measure the total system response. Short-term fluctuations in the array response can be monitored using the internal calibration subsystem. This is possible since the transmission properties through the telescope should be relatively constant.
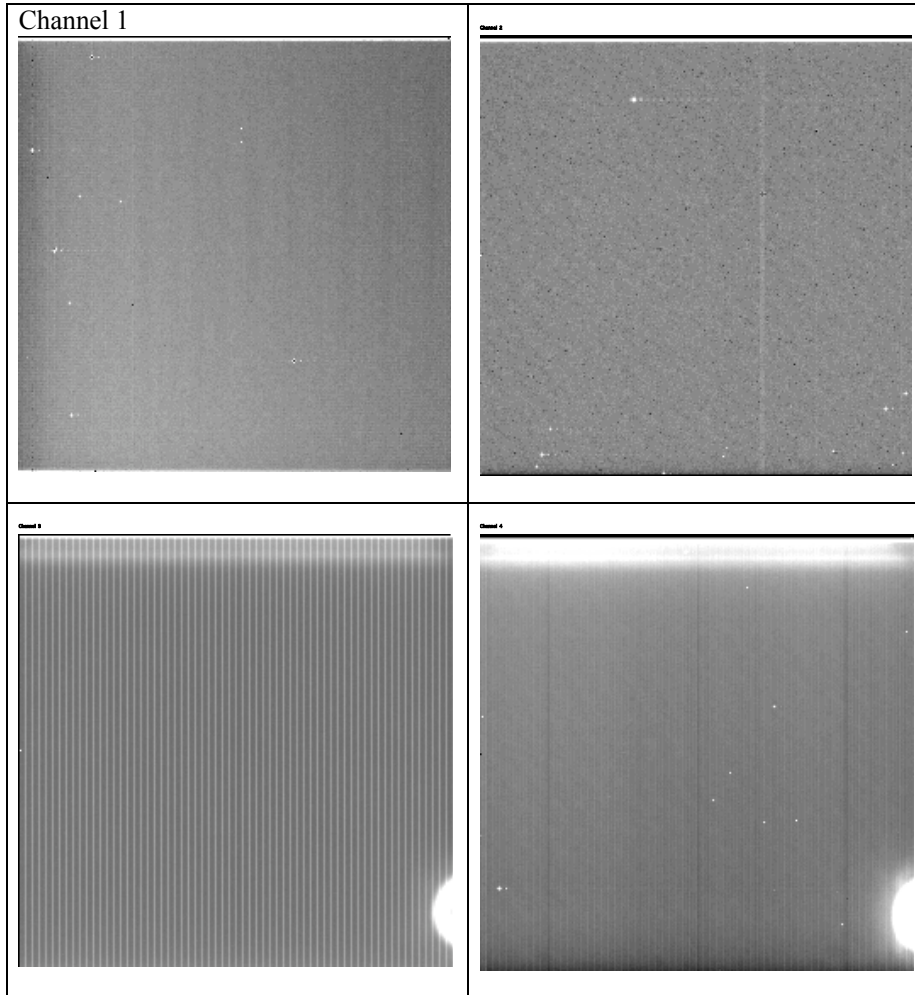
Figure 6.14: IRAC lab "dark" images. The bright regions near the lower right corner of the Si:As arrays (channels 3 and 4) are due to electronics near the edge of the array. Pixel (1,1), in the lower left corner, is the first pixel read out by the electronics

6.1.3.2.9   Effects of cosmic rays

Cosmic rays will strike the IRAC detectors fairly regularly (about one per array per second, each one affecting on average about 4 pixels), so strategies to deal with them in data reduction will have to be implemented. The most common cosmic rays do not affect the pixel performance in subsequent frames, and are confined to a few pixels around the peak. However, other events can cause

streaks or other multiple pixel structures in the array, and less common energetic events can cause muxbleed and residual images in subsequent frames. Single cosmic rays should not permanently affect the responsivity of a pixel, but multiple strikes over time, and exposure to strong solar proton events, may lead to an increase in the dark current of certain pixels, and perhaps to a degradation of responsivity. Despite significant solar events during normal operations, no significant degradation of the IRAC arrays has been observed.

Figure 6.15 shows portions of two 100-sec IRAC frames obtained during IOC/SV. Stars have been subtracted to show only cosmic rays in the images. The cosmic rays in channel 1 are more compact than those in channel 3, where cosmic ray effects can appear as streaks and blobs depending on the angle of incidence of the cosmic ray as it penetrates the thick Si:As detectors.
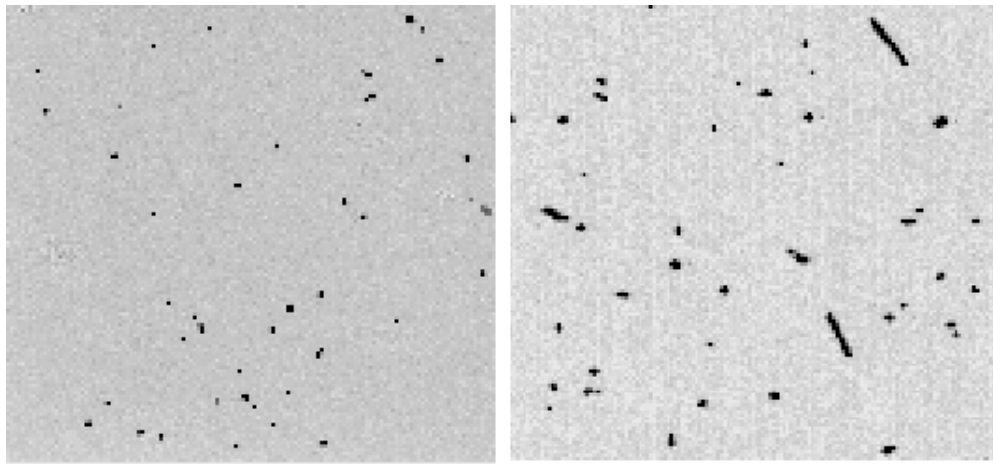
Figure 6.15: Channel 1 (left) and channel 3 (right) IRAC images (100 sec frame time) obtained during IOC/SV. Stars have been subtracted. All the dark pixels are affected by cosmic rays.
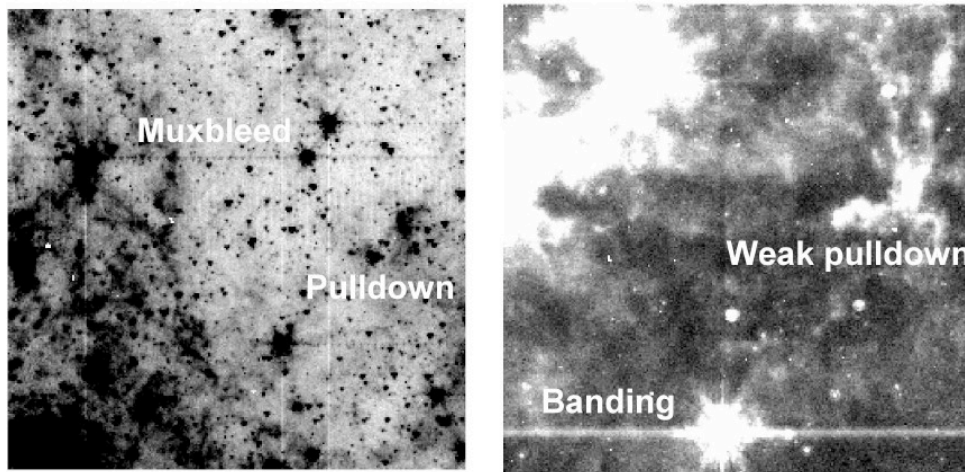
Figure 6.16: Images showing examples of image artifacts in channels 1 and 2 (left) and channels 3 and 4 (right). These images are early release observations of 30 Doradus taken during the IOC/SV period.

6.1.3.2.10   Thresholds for triggering image artifacts

Based on examination of 12 sec HDR data of various star-forming regions, it appears that several of the electronic artifacts seen in IRAC images (see Figure 6.16) are triggered by pixels above a certain fluence. These levels are likely to be appropriate for all frame times and should be sufficient for planning purposes. Table 6.5 provides estimates of how bright a pixel has to be to trigger a significant electronic artifact. Thresholds are in units of fluence (MJy/sr · sec). To find the threshold for a particular artifact for a given frametime, divide the appropriate value by the frametime. If the surface brightness of a pixel is above a threshold, that pixel will produce a specified artifact. To see if a point source may induce an artifact, convert the estimated flux density to surface brightness (assuming that the source is centered on an IRAC pixel) by scaling the flux density, in μJy, by 0.12, 0.12, 0.08, and 0.06 for channels 1–4, respectively. Keep in mind that, if a source is not centered on a pixel, its flux will be distributed more uniformly and be less likely to induce an artifact. The upper limit of the flux falling into the central pixel can be roughly estimated by multiplying the flux by the percent of flux in the central pixel (from Table 6.1) and the projected area of an IRAC pixel (~3.5×10^{-11} sr). Muxbleed affects the BCDs from left to right and will wrap around image rows, possibly affecting rows below the triggering pixel. Muxstriping is a stripe introduced in part of an image below a bright source, which may or may not continue to the bottom of a BCD. The bandwidth effect is a nonlinear residual in the readout electronics which produces a short, bright trail spaced 4, 8, and up to 12 pixels from a very bright triggering pixel.

Table 6.5: Fluence per second (in Mjy·sec/sr) of a pixel needed to trigger artifacts.

Ch	no_art	pull_dwn	mux_ovr	mux_str	mux_mod	mux_wk	banding	bandwth
1	9890.9	4200.0	6593.4	4945.5	1978.2	989.1	—	—
2	14965.0	4115.4	11223.7	9353.1	5986.0	2057.7	—	—
3	62652.6	4800.0	—	—	—	—	28000.0	70000.0
4	21273.7	10636.8	—	—	—	—	2925.1	1010.5

The definitions of the columns in Table 6.5 are

no_art: Pixels above this value will not produce artifacts in a BCD, the artifacts essentially saturate in the pedestal read and are canceled by the Fowler sampling. This value is above the saturation level of all detectors.

Pull_dwn: Threshold for column pull-down (or pull-up in channels 3 and 4).

Mux_ovr: Threshold above which the muxbleed artifact saturates and will not be visible. This value is above the saturation level of channels 1 and 2.

Mux_str: Threshold for producing a strong instance of muxbleed. Strong muxbleed results in a noticeable trail of affected pixels up to 128 samples away for channel 1 and 192 samples away for channel 2. Note that muxbleed affects every fourth pixel along the readout direction of the InSb arrays.

mux_mod: Threshold for producing a noticeable trail up to 64 samples away from the triggering pixel for both channels 1 and 2.

mux_wk: Threshold for producing a noticeable muxbleed trail up to 32 samples away from the triggering pixel for channels 1 and 2.

banding: Threshold above which horizontal banding occurs for channels 3 and 4. As some of the banding is optical in nature, banding can occur at lower thresholds if more than one pixel in the row contributes.

bandwth: Threshold above which the bandwidth effect (limited version of muxbleed affecting the fourth, eighth, and 12th sample after the triggering pixel) occurs for channels 3 and 4.

For more information on IRAC instrumental artifacts, please see the IRAC Data Handbook.

6.1.3.2.11 First-frame effect

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 on each of the four outputs of an array (pinstriping) and a small spatial gradient across the image. 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 and temperature changes. 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 image. The largest dark offsets occur when a frame is taken with a very short interval from the preceding image, which occurs when multiple frames are commanded at once (using "repeats"). In lab testing, this effect came to be known as the "first frame effect" for this reason. We eliminated commanded repeats from IRAC science observations, so there is always a finite interval between one frame and its predecessor. Most observers should be dithering between frames, meaning that the interval between frames will be in the range 8 to 60 sec. The first frame of every AOR with frame time ≥12 sec is taken with the HDR mode (see Table 6.9) to mitigate the first-frame effect.

The "first frame effect" was calibrated in ground tests with the shutter closed. Figure 6.17 shows the dependence of dark offset on the interval between 30 sec 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 sec frames (see

Table **6.7**). Channel 1 has very significant variations in offset, compared to the low background in that band. 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 (note the scale in Figure 6.17).
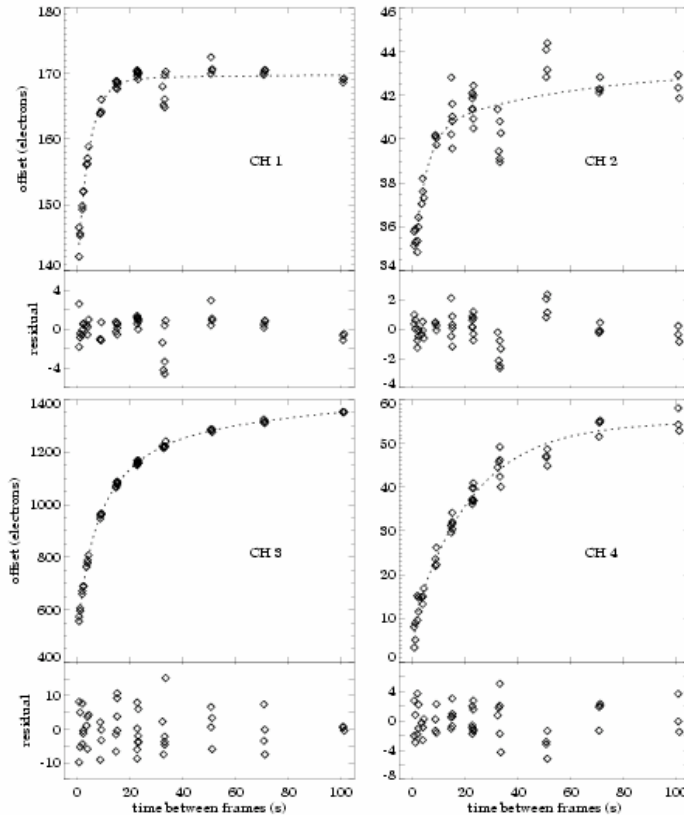
Figure 6.17: First-frame effect – electronic offset as a function of time interval between frames.

6.1.3.2.12 *Stability*

In-flight data indicate that IRAC is very stable. The gain is monitored during campaigns using a network of calibration stars, and the offset is monitored by observing a relatively dark and empty patch of sky in the northern continuous viewing zone. Primary calibration stars are observed at the beginning and end of each campaign, allowing an assessment of the very-long-term stability. The rms of the flux measurements across the first 12 campaigns (Dec 2003 through Aug 2004) for individual calibration stars is less than 2% at all wavelengths. We are searching for subtle trends and correlations with observatory events, but none have been found. Secondary calibrators are observed every 12 hr during each campaign, allowing an assessment of the intra-campaign stability. Figure 6.18

shows an example of the histogram for observations of HD 113745, one of the secondary calibrators observed during the 10th IRAC campaign. The rms variation in the measured source flux is 0.9, 0.9, 0.7, and 1.5% in channels 1, 2, 3, and 4, respectively. Other stars yield similar results. Care must be taken when comparing photometry from one observation to another; in particular, observations on different parts of the array need to have a corrective flat-field applied, in order to derive accurate relative photometry. (This was done for the figure shown here, which combines 5 locations widely spaced on the array. Even without this correction, observations of stars put in the same part of the array would yield comparable [slightly better, actually] repeatability.)
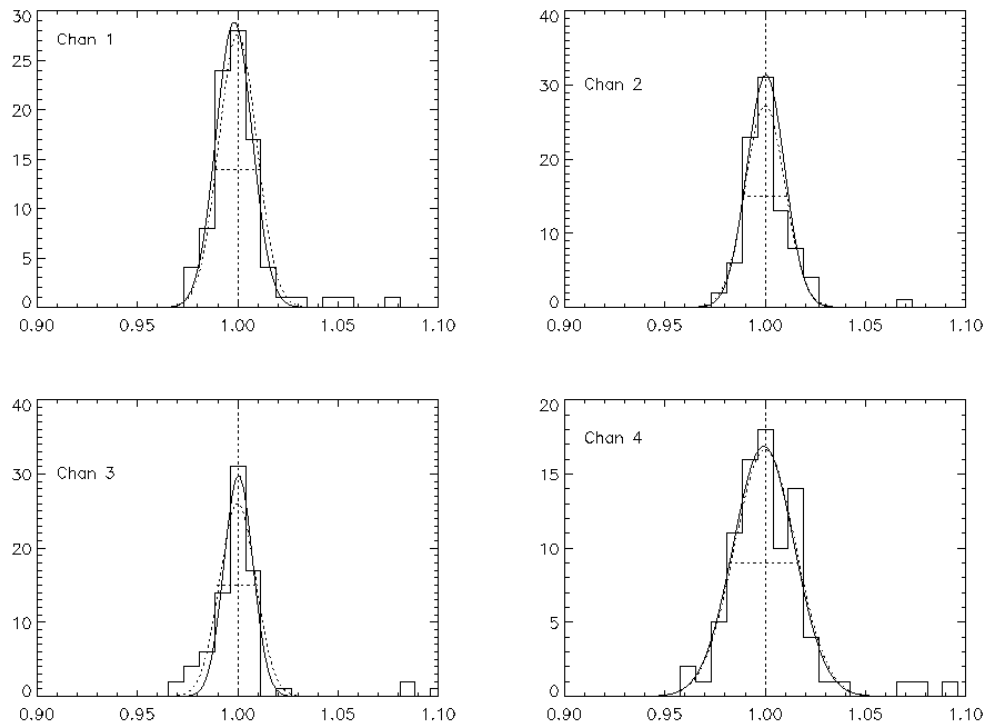
Figure 6.18: Histograms of photometric observations of HD 113745, one of the IRAC secondary calibrators, as observed during the 10[th] IRAC campaign. Each panel shows one IRAC channel (1: top left; 2: top right; 3: bottom left; 4: bottom right). In each panel, the horizontal axis is the flux measured from an individual image, divided by the median flux for all images, and the vertical axis is the number of observations per bin. The solid curve is a Gaussian fit to the histogram. The dotted lines show the Gaussian distribution predicted for the measured rms of the distribution. Apart from a few positive outliers, which are most likely due to cosmic rays, the distributions appear normal.

The absolute offset in the IRAC images also appears stable. There is a shutter within IRAC that is not being operated in flight. Therefore, we refer to ground tests for the absolute stability and in-flight tests for verification. The main IRAC stability test was performed at Ball Aerospace, when the cryogenic telescope assembly was isolated in a quiet chamber and put into superfluid conditions. During the 3-day stability test, a long set of relatively flight-like command sequences was executed on IRAC, to determine whether it would be possible to

use IRAC in space without the shutter. Figure 6.19 shows the DC offset and rms measured during the 3-day test for the set of pseudo-observations mimicking the Great Observatories Origins Deep Survey (GOODS) observing strategy, which uses 200 sec frames. Seven pseudo-observations were performed, spread throughout the 3-day test. Circles (red) show measurements for the central 51 × 51 pixels of the images, while small dots (black) show the same measurements for subsections separated into each of the four readout amplifiers. The deviant rms in channel 1, frame 40 is likely due to a cosmic ray. Several instabilities are evident in the data. First, it is easy to see the seven separate observations have different offset levels. This is most clearly evident in channel 1, and is evident but very small in channel 2. Second, the first frame of each AOR is discrepant. This so-called "first-frame effect" is described in Section 6.1.3.2.11 and in the IRAC Data Handbook. It can be partially corrected by the pipeline. Third, there is a gradual, monotonic trend within each AOR. The trends are at the 2 DN level for channels 1, 2, and 4, and at the 10 DN level for channel 3. These trends are partly due to the "first-frame effect" (Section 6.1.3.2.11) and partly due to latent response (Section 6.1.3.2.3) from bright lamps used in the preceding activities. In all cases, the trends are much smaller than the pixel-to-pixel sky noise. For channel 1, the instabilities in the offset are larger than the sky brightness for very short frames, mostly because the celestial background is very faint at that wavelength, but the offset instability is less than 10% of the celestial background for 200-sec frames. In channels 3 and 4, the sky brightness is much higher, and offset instabilities are less than 1% of the sky brightness.

6.1.4 Electronics

6.1.4.1 Hardware

IRAC has no moving parts (other than the shutter, which is not operated in flight). The instrument takes data by staring at the sky and sampling the arrays between resets. IRAC is capable of operating each of its four arrays independently and/or simultaneously. All four arrays are used during normal, full-array operation.

6.1.4.1.1   Fowler sampling

Multiple (Fowler) sampling is used to reduce the effective read noise. This mode of sampling consists of taking N non-destructive reads immediately after the reset, and another N non-destructive reads near the end of the integration. Differencing is performed in the IRAC electronics to generate one integer value per pixel per exposure to store on the spacecraft and transmit to the ground. The Fowler N used for an observation depends on integration time and was selected to maximize the S/N, based on in-flight performance tests.

6.1.4.1.2   Exposure time and frame time

The relationship between the exposure time (T_{ex}) and frame time is shown in Figure 6.20. The exposure time is defined as the time elapsed between the first pedestal sample and the first signal sample. The Fowler samples are taken consecutively at 0.2-second intervals (0.01 seconds for subarray) in each group

(pedestal and signal samples). The frame time (T_f - T_i) is the total time elapsed between resets which could include multiple reads and dead time before and after Fowler sampling. The frames are commanded by specifying the number of Fowler samples for the pedestal (N_F) and the number of "wait ticks" in between the pedestal and signal frames (N_W); then the frame time is $T_F = (2N_F + N_W)\tau$, where τ is the readout time (0.2 sec for full array, 0.01 sec for sub-array mode).
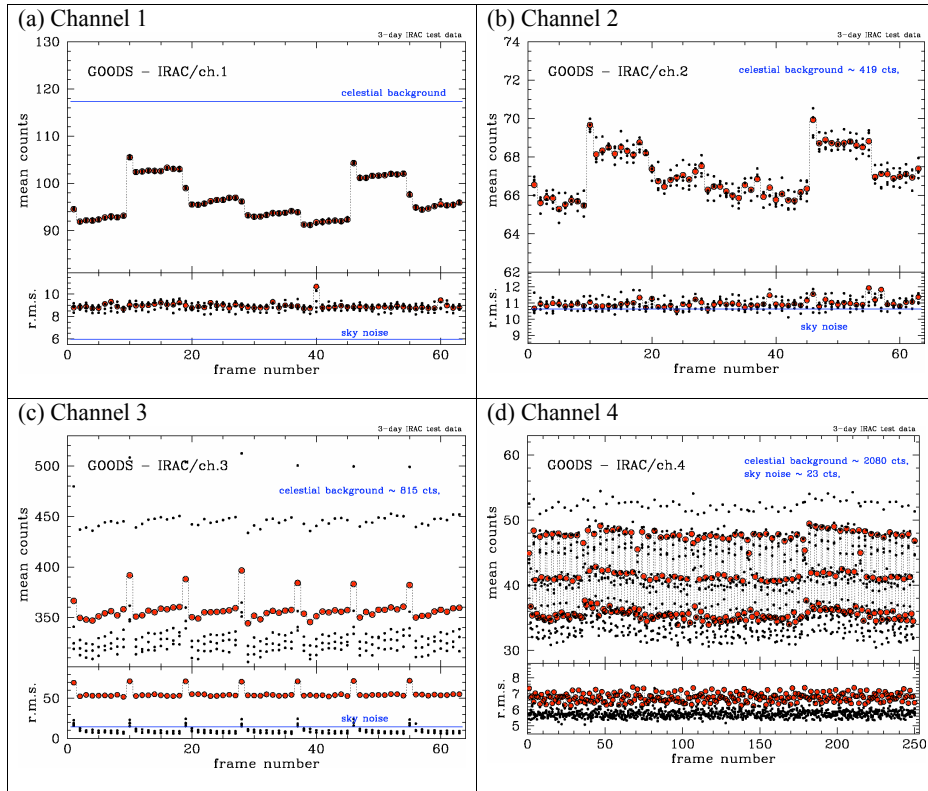
Figure 6.19: IRAC pre-launch laboratory stability measurements using 200-second frames. Channel 4 uses 4 × 50 sec exposures, instead of 200 sec, hence the multiple sets of datapoints in the channel 4 plot.

6.1.4.1.3   Subarray mode

In subarray mode, only one corner, 32×32 pixels offset by 8 pixels from the edges, is read out from one array. Pixels (9–40, 9–40) of the array are read out. The subarray pixel size is the same as the full array pixel size (~1.2″). Fowler sampling is performed as in full array mode, but a set of 64 subarray images are generated and tiled into a single 256×256 image before data are sent from IRAC. In subarray mode, Fowler sampling is performed at 0.01 sec intervals. Subarray mode is useful for observing very bright sources and for obtaining high temporal resolution.

6.1.4.1.4   Calibration Lamps

IRAC contains two types of internal calibration lamps. The transmission
calibrator lamps are designed to illuminate all four arrays and provide an internal
responsivity measurement. There are two transmission calibrator spheres, each of
which contains two lamp elements. To illuminate the arrays, the shutter is closed,
a transmission lamp is turned on, and the light from that lamp bounces off a
mirror on the back of the shutter. The flood calibrators individually illuminate
each detector. The flood calibrators can be controlled individually, and they can
be used whether the shutter is open or closed. The flood calibrators are operated
at the end of each IRAC campaign and used as a consistency check. Calibration
of IRAC observations is described in section 6.3.1.
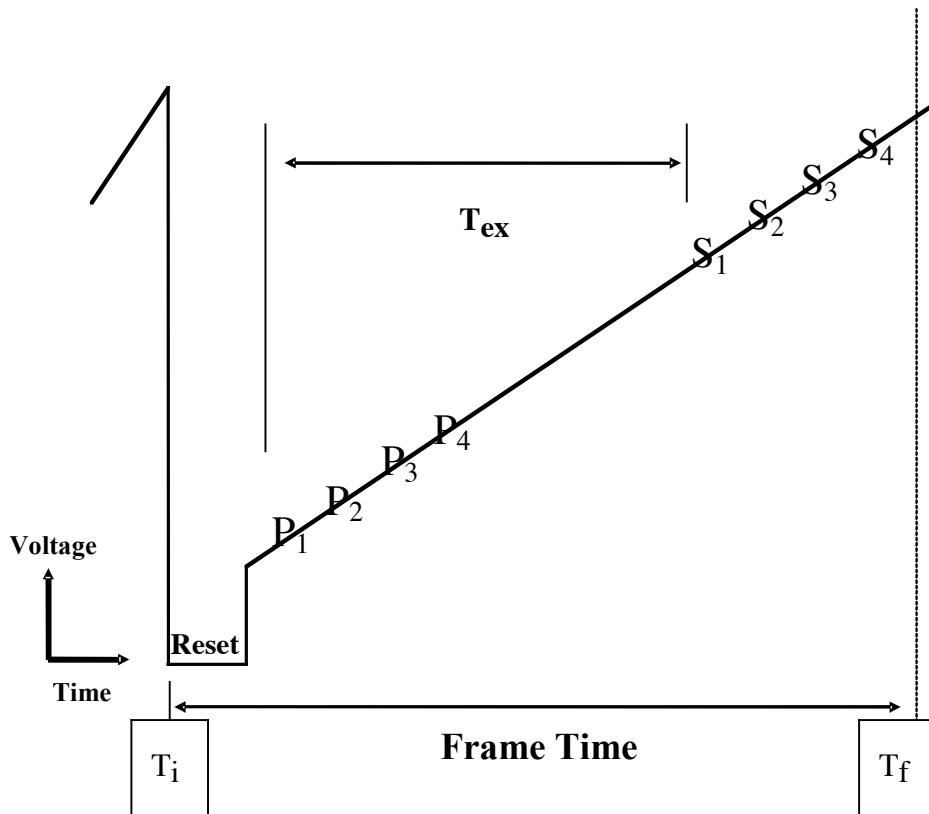
**Figure 6.20: Fowler sampling times for one pixel (Fowler *N*=4). The *P_n* (*n*=1,2,3,4) show the
"Pedestal" readouts, and the *S_n* show the "Signal" readouts. *T_{ex}* is the effective exposure
time, and *T_f − T_i* is the "frame time," or total time to obtain one IRAC image. The reset part
of the sketch is not at the same time and voltage scale as the rest of the figure.**

6.1.4.2 Firmware

The IRAC firmware controls the focal plane assemblies, calibration electronics,
and warm electronics boards. Apart from autonomous fault protection, the IRAC
firmware responds only to commands sent by the Spitzer Command & Data
Handling (C&DH) computer. The C&DH sends setup commands to configure the
electronics, requests for each telemetry packet, and integration commands to
generate images. IRAC responds to each command with an acknowledgment. In

the case of a command that requests telemetry, the acknowledgment consists of the telemetry packet, which is sent on the low-speed connection between IRAC and the C&DH. There are two types of engineering data: special engineering data, which are collected every 4 sec, and housekeeping data, which are collected every 30 seconds. Special engineering data are used for onboard communication between IRAC and the C&DH, while housekeeping data are used on the ground to monitor instrument performance. A command that generates images from the arrays is acknowledged on the low-speed line, and when the frame is complete, the "data-ready" signal is sent on the high-speed line. The frames (together with their ancillary data) are then transferred one at a time to the C&DH. The rate of transfer is 2 seconds per frame, which limits the data collection rate to a maximum of four frames every 8 seconds for IRAC observations with all four arrays.

Autonomous fault protection ensures that none of the monitored voltages or currents enters into a red limit. Fault protection is performed by a "watchdog" demon that is always running when the instrument is on. The red limits are stored in IRAC memory. If a voltage in the focal plane assembly goes into a red limit, IRAC will turn off the affected focal plane array. (The individual pixel values are not monitored, so a bright astronomical source should not trigger a red limit.) The command sequences will continue to execute; therefore, it is possible for normal completion of an IRAC observing campaign to occur with only three of the four arrays returning data. If a second focal plane assembly has a telemetry datum go into a red limit, then IRAC will send the C&DH (via the special engineering data on the low-speed line) a request to be turned off. If a telemetry point other than one affecting a single focal plane goes into a red limit, IRAC will also send the C&DH a request to be turned off.

6.2 How to Use IRAC

6.2.1 Performance of the Instrument

To estimate the sensitivity of IRAC in flight, where possible we use the measured properties of Spitzer and IRAC from IOC/SV; otherwise we use the required performance based on the design specifications. The online tool, SENS-PET (at http://ssc.spitzer.caltech.edu/tools/senspet/), makes this estimation for you and should be used for planning observations. The sensitivity to point sources (in flux density units) is based on the following formula:

$$\sigma = \frac{\sqrt{N_{pix}}}{ST_{ex}f_p}\sqrt{BT_{ex} + (BT_{ex}f_F)^2 + R^2 + DT_{ex}}, \qquad (6.1)$$

where the scale factor is

$$S = \frac{Q\eta_T\eta_I A\Delta\lambda}{h\lambda}, \qquad (6.2)$$

the background current is

$$B = SI_{bg}f_S\Omega_{pix}f_{ex}, \qquad (6.3)$$

and the effective exposure time is

$$T_{ex} = T_F - 0.2N_F . \qquad (6.4)$$

In these equations,

- the spectral resolving power $\lambda / \Delta\lambda$ is from Table 6.3;
- the detector quantum efficiency Q (electrons per photon) is from Table 6.4;
- the instrumental throughput η_I is from Table 6.3;
- the telescope throughput $\eta_T =$[0.889, 0.902, 0.908, 0.914] for channels 1 to 4, respectively (with Be primary, Al-coated secondary, and 50 nm ice contamination);
- the telescope area (including obstruction) A=4636 cm^2;
- the equivalent number of noise pixels N_{pix} is from Table 6.1 (and defined in section 6.1.2.2);
- h is the Planck constant;
- I_{bg} is the background surface brightness in MJy/sr;
- f_S=1.2 is the stray light contribution to the background;
- the dark current D is given in section 6.1.3.2.7;
- the read noise R is from Table 6.4;
- Ω_{pix} is the pixel solid angle (see Table 6.1);
- f_p is the in-flight estimated throughput correction for point sources (Table 6.6);
- f_{ex} is the in-flight estimated throughput correction for the background (Table 6.6).

The "throughput corrections" f_p and f_{ex} were determined by comparing observed to expected brightness of stars and zodiacal light. Stars were measured in a 10-pixel radius aperture, and the zodiacal light was measured in channels 3 and 4 for comparison to the COBE/DIRBE zodiacal light model. (This measurement is not possible in channels 1 and 2 because we cannot use the shutter for absolute reference.)

Early in the mission, we found that the throughput in channels 3 and 4 was lower than expected, both for extended emission and point sources (but more so for point sources). Measurements of diffuse Galactic emission confirmed the deficit, and measurements of the PRF using bright stars showed that a considerable amount of stellar flux was being spread all across the arrays. The deficit in throughput from diffuse emission was due to the fact that the QE of the arrays had been overestimated. No reliable measurements of QE existed for channels 3 and 4, so the QE was based on a theoretical model that incorrectly assumed that all the flux was reflected at the front of the detector diode chip (the detector arrays are backlit). Some of the light that passes through the detector is scattered widely across the array. Measurements on a sister array confirmed this internal scattering, and showed that it is strongly wavelength-dependent. Further laboratory studies

are underway. There is considerable evidence that the "loss" of QE and the scattered light are not due to contamination, or damaged optical coatings, etc.

Table 6.6 lists some useful combinations of IRAC instrument parameters.

Table 6.7 gives the background brightnesses, in useful units, for three nominal observing directions. The low-background model applies near the ecliptic pole; the high-background case is in the ecliptic plane; and the medium-background case is intermediate. The background model includes contributions from emission and scattering from zodiacal dust and emission from Galactic dust. The near-infrared cosmic infrared background radiation is not included because it will be partially resolved by Spitzer.

Table 6.6: Useful quantities for IRAC sensitivity calculations

Wavelength	3.6 μm	4.5 μm	5.8 μm	8 μm
Conversion factor (electrons/sec)/(MJy/sr)	25	29	14	29
S (electrons/sec)/(μJy)	0.77	0.89	0.42	0.91
Gain (electrons/DN)	3.3	3.7	3.8	3.8
f_p (throughput correction for point sources)	1.06	0.84	0.45	0.61
f_{ex} (throughput correction for background)	1	1	0.72	0.88
Zero magnitude flux densities (Jy)	280.9	179.7	115.6	64.13

Table 6.7: Background brightness in IRAC wavebands

	3.6 μm	4.5 μm	5.8 μm	8 μm
"low" background model				
$I_v f_S$ (MJy/sr)	0.093	0.32	1.7	6.6
F_v^{BG} (μJy)	3.2	11	57	220
B (elec/sec)	2.5	9.9	18	184
"medium" background model				
$I_v f_S$ (MJy/sr)	0.15	0.44	2.3	9.3
F_v^{BG} (μJy)	5.1	15	79	320
B (elec/sec)	4.1	14	25	260
"high" background model				
$I_v f_S$ (MJy/sr)	0.52	1.0	5.6	22
F_v^{BG} (μJy)	18	35	190	750
B (elec/sec)	14	32	60	620

The quantity f_F is the flat field pixel-to-pixel variance, which depends on the observing strategy. In what follows, we will set $f_F = 0$, which would apply strictly in the case of stable detectors with perfect flat field measurements, and should apply practically for highly-dithered observations. An observation with no dithering will be limited by the correlated noise. The accuracy of a flat field derived from a single observing campaign was measured to be 2.4, 1.2, 1.0, and

0.3% in channels 1, 2, 3, and 4, respectively, by comparing flats in several campaigns. Using combined flats ("super sky flat") from the first two years, the estimated f_F is 0.4, 0.2, 0.2, and 0.05% in channels 1–4, respectively. Using these values for f_F in equation 6.1, single frames are dominated by background and read noise. When combining multiple frames to generate a mosaic, the background and read noises will average down (as square root of the number of frames), while the flat-field noise will only average down for dithered observations. For N undithered observations on the "medium" background, flat-field noise dominates when the total exposure time, $N \times T_{ex}$, exceeds approximately 420 sec (using individual campaign flats) or 2.5 hrs (using the super sky flat). For dithered observations, the flat-field noise will also average down, and will only be important for the very deep observations of high background fields.

For the frame times used in IRAC operations in flight, Table 6.8 gives the readout mode and Fowler number. For full array readout mode, only the 2, 12, 30, and 100 sec frame times can be chosen in the IRAC AOT; the 0.6 and 1.2 sec frame times come as part of the "high dynamic range" (HDR) sequences. The 0.4 sec full frame time is only available for channels 1 and 2 in Stellar Mode. The frame sets that are taken for each pointing in HDR mode are shown in Table 6.9. Long frame times at 8 μm are background-limited. Therefore there is a maximum frame time of 50 sec at 8 μm, and the "100 sec" frames are automatically converted into two repeats of 50 sec frames. The last column, T_h, gives the extra time spent taking the HDR frames (used in the observing time estimate equation below).

Table 6.8: Fowler numbers for IRAC frames

Frame Time (sec)	Readout Mode	Fowler Number	Wait Ticks
100	Full	16	468
50	Full (8 μm)	16	218
30	Full	16	118
12	Full	8	44
2	Full	4	2
1.2	HDR	1	4
0.6	HDR	1	1
0.4	Stellar	1	0
0.4	Subarray	8	24
0.1	Subarray	2	6
0.02	Subarray	1	0

Table 6.9: IRAC High-Dynamic-Range (HDR) framesets

Long Frame Time	List of frames taken	T_h (sec)
100	0.6, 12, 100	15
30	1.2, 30	3
12	0.6, 12	2

Table 6.10: IRAC point-source sensitivity, low background (1σ, μJy)

Frame Time (sec)	3.6 μm	4.5 μm	5.8 μm	8 μm
100	0.60	1.2	8.0	9.8
30	1.4	2.4	16	18
12	3.3	4.8	27	29
2	32	38	150	92
0.6[a]	180	210	630	250
0.4[b]	360	430	1260	450
0.4[c]	81	89	609	225
0.1[c]	485	550	2010	690
0.02[c]	7300	8600	25000	8100

Table 6.11: IRAC point-source sensitivity, medium background (1σ, μJy)

Frame Time (sec)	3.6 μm	4.5 μm	5.8 μm	8 μm
100	0.73	1.4	9.3	12
30	1.6	2.8	18	21
12	3.6	5.3	31	34
2	32	38	150	110
0.6[a]	180	210	640	260
0.4[b]	360	430	1260	460
0.4[c]	82	89	610	250
0.1[c]	490	550	2020	720
0.02[c]	7300	8600	25000	8100

Table 6.12: IRAC point-source sensitivity, high background (1σ, μJy)

Frame Time (sec)	3.6 μm	4.5 μm	5.8 μm	8 μm
100	1.3	2.1	14	18
30	2.5	4.1	27	32
12	4.8	7.1	44	52
2	34	41	180	156
0.6[a]	180	220	660	330
0.4[b]	360	430	1280	540
0.4[c]	84	93	650	340
0.1[c]	490	560	2100	860
0.02[c]	7300	8600	25000	8200

[a] available only in high-dynamic-range mode.
[b] available only in stellar photometry mode.
[c] subarray mode (set of 64 32×32 images). Sensitivity is per frame, not the sensitivity of a 64-frame coadd.

Table 6.10 to Table 6.12 and Figure 6.21 to Figure 6.23 show the sensitivities for the four IRAC channels for each of the three background models. The sensitivities in the tables are for point sources extracted from single images (but perfectly flat-fielded). In the figures, the sensitivities are for point sources extracted from coadded images (perfectly registered). We do not include

"confusion noise" (due to overlapping images of distant galaxies or other sources of background structure) in the sensitivity estimates. The detectors are assumed to perform according to the IRAC detector measurements of read noise, dark current, and quantum efficiency. The first 6 rows in each table show the sensitivity for full-array readouts, and the last three rows show the sensitivity for subarray readouts.
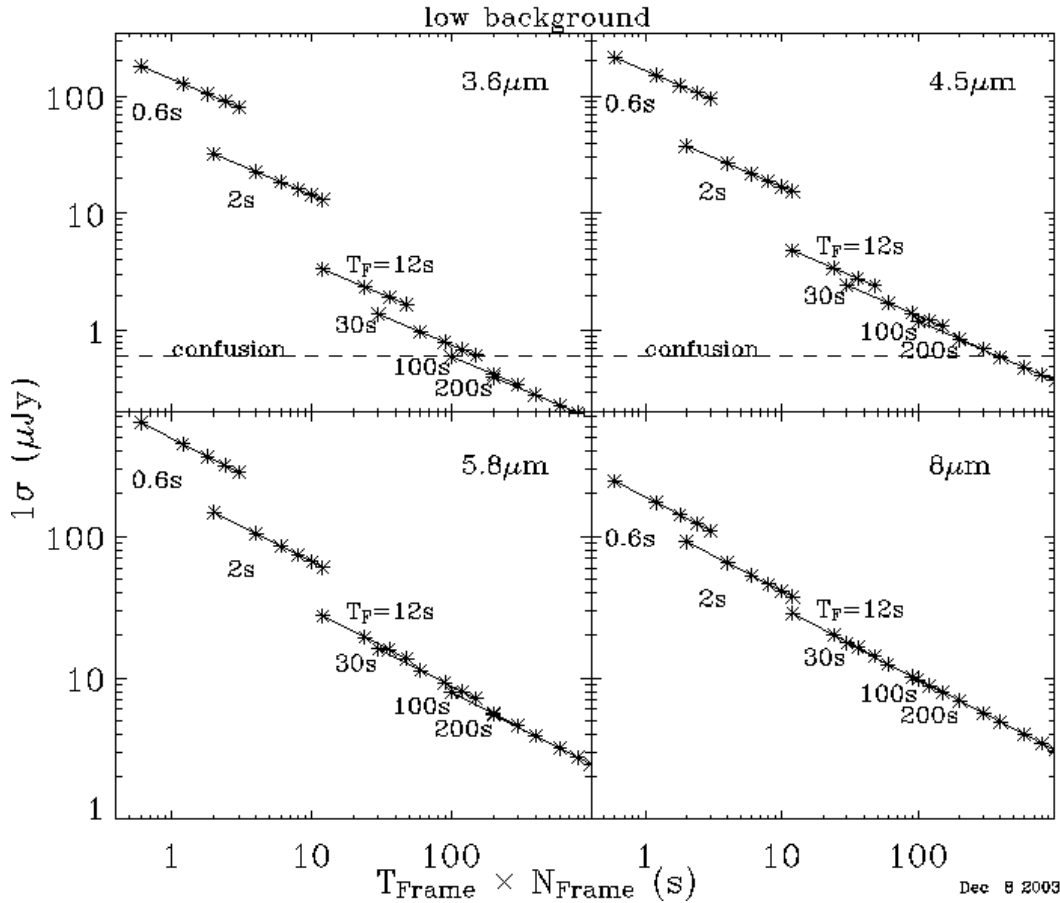
Figure 6.21: IRAC point source sensitivity as a function of frame time, for low background. To convert to MJy/sr, see formula 6.8.

Figure 6.21, Figure 6.22, and Figure 6.23 show the point source sensitivity as a function of integration time for each background model. The "time" axes in the plots represent the frame time for the images, which does not include time for moving the telescope. The IRAC full-array frame times are 0.6, 2, 12, 30, and 100 seconds. Other times plotted below are assumed to use multiple exposures of those fixed times.

For bright sources, shot noise due to counting statistics in electrons from the source itself becomes the dominant source of noise. We can estimate the total noise by adding the shot noise in quadrature, so that

$$\sigma_{tot} = \sigma\sqrt{(1+(F/F_b))} \qquad (6.5)$$

where σ is the noise from equation 6.1 and

$$F_b = ST_{ex}\sigma^2 . \qquad (6.6)$$

In the bright source limit, $F \gg F_b$, then the signal-to-noise ratio becomes

$$S/N = \sqrt{ST_{ex}F} . \qquad (6.7)$$

If the exposure time is in seconds and the source flux is in μJy, then for IRAC channels 1, 2, 3, and 4, respectively, *S/N* is 0.88, 0.95, 0.65, and 0.95 times $\sqrt{T_{ex}F}$. Thus, for example, to design an experiment that requires a *S/N*=1000, for a star with a brightness of 1 mJy at 3.6 μm, we require an exposure time of 1290 sec. This could be achieved, for example, using 13 frames of 100 sec, because the star is below the saturation limit for 100 sec frames.
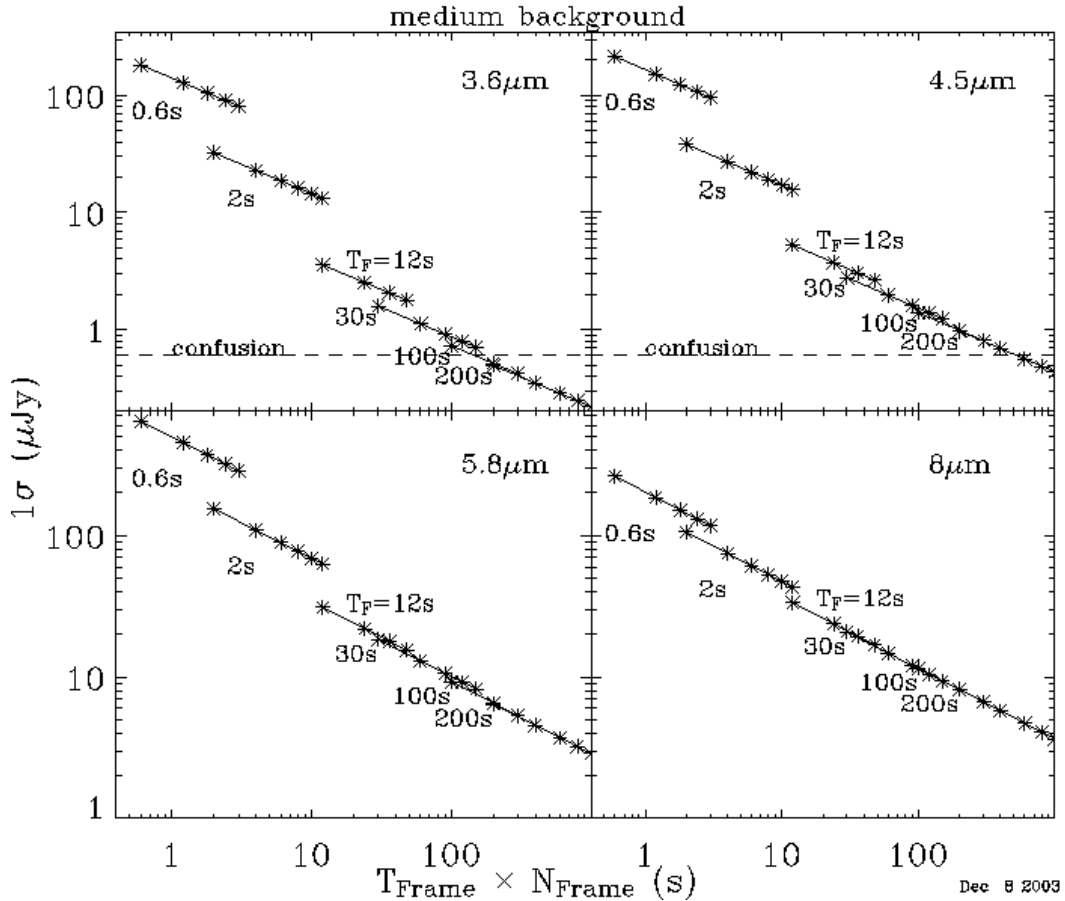
Figure 6.22: IRAC point source sensitivity as a function of frame time, for medium background. To convert to MJy/sr, see formula 6.8.

In the sensitivity figures, the dashed line at 0.6 μJy is the confusion limit predicted by Franceschini et al. (1991, A&AS, 89, 285). This does not represent a hard sensitivity limit, but rather indicates where source confusion affects

reliability of source extractions for low background regions. Data from IOC/SV show noise decreasing as $\sqrt{N}$ to 0.25 μJy (channels 1 and 2) or 0.6 μJy (channels 3 and 4). Moderately deep source counts indicate that a source density equivalent to 36 beams/source is reached at 20.5 mag, or 1.8 and 1.1 microJy at 3.6 and 4.5 μm, respectively (Fazio et al. 2004, ApJS, 154, 10). The confusion estimates by Franceschini et al. and Fazio et al. are for low background, extragalactic observations only. For observations of higher background or more "cluttered" regions (such as the Galactic Plane) the confusion noise will be much more significant. We recommend that observers examine the archive for available public data from regions similar to their targets to determine appropriate measures of confusion noise.
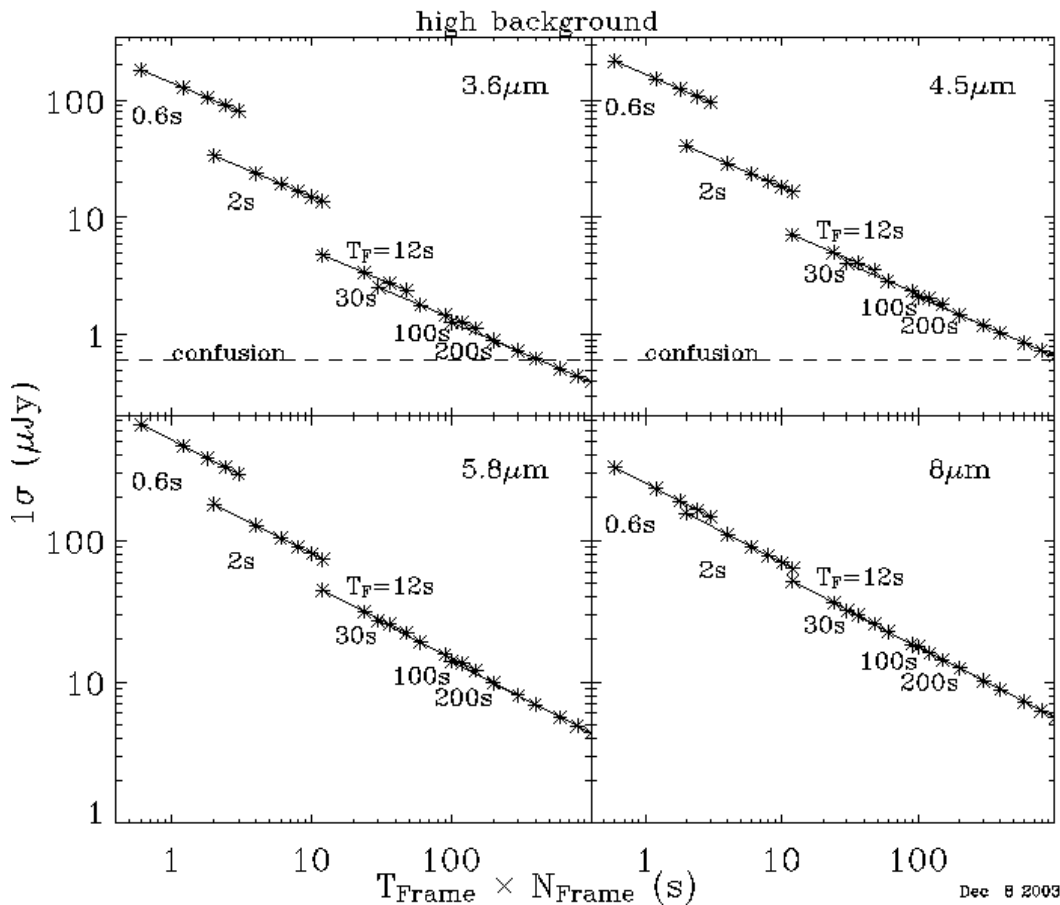
Figure 6.23: IRAC point source sensitivity as a function of frame time, for high background. To convert to MJy/sr, see formula 6.8.

For diffuse emission, the surface brightness sensitivity per pixel (in MJy/sr) is

$$\frac{0.03 f_p}{f_{ex}\sqrt{N_{pix}}} \times \text{ the point source sensitivity [in μJy]}. \tag{6.8}$$

(The noise pixels, N_{pix}, are defined in section 6.1.2.2.)

The saturation limit for IRAC is calculated as follows. Using the same notation as earlier in this section,

$$F_{sat} = \frac{Wf_w - BT_F f_{ex}}{ST_F f_{cen} f_p} \tag{6.9}$$

where W is the well depth (Table 6.4), $f_W = 0.9$ is the fraction of the well depth to which we can linearize the intensities, and f_{cen} is the fraction of the source flux falling onto the central pixel (Table 6.1). Table 6.13 shows the point source saturation limits of IRAC at each frame time. If you are observing an extremely bright piece of sky, such as an H II region, then the saturation limit is lower; you should include your sky brightness in the background B; for the saturation table, we used the "medium" diffuse sky background from

Table **6.7**. Note that the saturation value is conservatively computed from the worst-case in which the PSF is directly centered on a pixel. While users may find that in their actual data the saturation value of point sources is higher, the values given in Table 6.13 should be used for planning purposes.

To apply Table 6.13 for extended sources,

$$I_{sat} = 28.6 \times f_{cen} F_{sat} \quad \text{for compact (diameter} < 30'') \text{ sources} \tag{6.10}$$

$$I_{sat} = 28.6 \times \frac{f_{cen} f_p}{f_{ex}} F_{sat} \quad \text{for more extended sources,}$$

where I_{sat}^{ext} is the total surface brightness (in MJy/sr) at which a pixel saturates; Ω_{pix} is the solid angle, in sr, subtended by the pixel; f_p, f_{cen}, and f_{ex} are as defined above, and F_{sat} is the saturating point source flux density for (in mJy) from Table 6.13 appropriate for the channel and integration time. For 8.0 μm observations at low ecliptic latitude, an estimate of zodiacal light should be included in the surface brightness.

Table 6.13: Maximum unsaturated point source (in mJy), as a function of IRAC frame time

Frame Time (sec)	3.6 um	4.5 μm	5.8 μm	8.0 μm
100	3.8	3.9	27	28
30	13	13	92	48
12	32	33	230	120
2	190	200	1400	740
0.6	630	650	4600	2500
0.4*	950	980	6950	3700
0.4**	1000	820	3100	2300
0.1**	4000	3300	13000	9000
0.02**	20000	17000	63000	45000

*stellar mode; **subarray mode

The zodiacal background only makes a difference for long frames in channel 4 when observing near the ecliptic plane. If the bright extended source extends well beyond the 5′×5′ FOV, then the saturation brightness is *lower* by the factor f_s.

Use Table 6.13, in conjunction with Table 6.10 through Table 6.12, to determine the dynamic range of an individual IRAC frame. If bright point sources (e.g., IRAS sources) are in your field of view, and you want to measure their fluxes as well as those of fainter sources, then you can use the "high dynamic range" (HDR) option. With this option, shorter frames will be taken in addition to your specified frame time.

6.2.2 Observing Time Estimation

The Spot tool provides fast and accurate estimates of observing time for IRAC. To optimize AORs, however, it may be useful to bear the following information in mind.

Each AOR duration includes a flat 215 sec slew overhead to the initial position (an additional 300 sec overhead is charged for moving targets). Command overheads between exposure repeats are very short (a few tenths of a second), except for the 2-sec full array and 0.02-sec subarray frame times, where a delay is put in to ensure that the data can be transferred from the IRAC buffer to the spacecraft (for which the maximum rate is 8 sec per 4 full array frames).

Slew durations and allowed settle times are a function of slew length. Slews between points in dither patterns, maps or cluster targets (apart from those between cluster targets specficied by celestial positions) are carried out under gyro control in low-rate mode. Each slew and settle has a command overhead of approximately 2 sec. Slews of 100" or less have a settle time of 4 sec; slews between 100" and 300", 5 sec; and, slews of more than 300", 6 sec. The slew times themselves are a function of the slew length. Slews from 0"-100" take 0-4.2 sec (excluding overhead), and those from 100"-300" take 4.2-8.4 sec. Slews of a degree take 54 sec. For cluster targets specified by celestial coordinates, a different slew mode is used, with larger command overheads, but with the ability to switch to high-rate mode for longer slews. Overall, slews of >0.5 deg are faster using this mode. AORs with widely spaced cluster targets will therefore be more efficient if RA, Dec positions are given for the cluster points.

Observing strategy is a series of compromises between sensitivity, redundant coverage, and wall-clock time. We recommend a given sky position be visited a minimum of three times (and preferably, five), in order to mitigate the effects of radiation hits and bad pixels. We also recommend that dithering be performed to mitigate the effects of pixel-to-pixel gain variations. Repeat observations at a given sky position should sample different parts of the array, either by using a medium or large dither scale, and/or by making approximately half-array offsets between map positions or between AORs covering the same part of the sky. This technique also ensures good straylight removal. For more details on recommended observing strategies, see section 6.2.4.

6.2.3 Astronomical Observation Template (AOT) Description

The IRAC AOT consists of an (optional) dither pattern superposed on an (optional) rectangular-grid raster. In this section, we describe the manner in which an IRAC AOR is specified by an observer, and the activities that the spacecraft will perform when the specified AOR is executed. Figure 6.24 shows an example of an IRAC AOT dialog box in Spot.
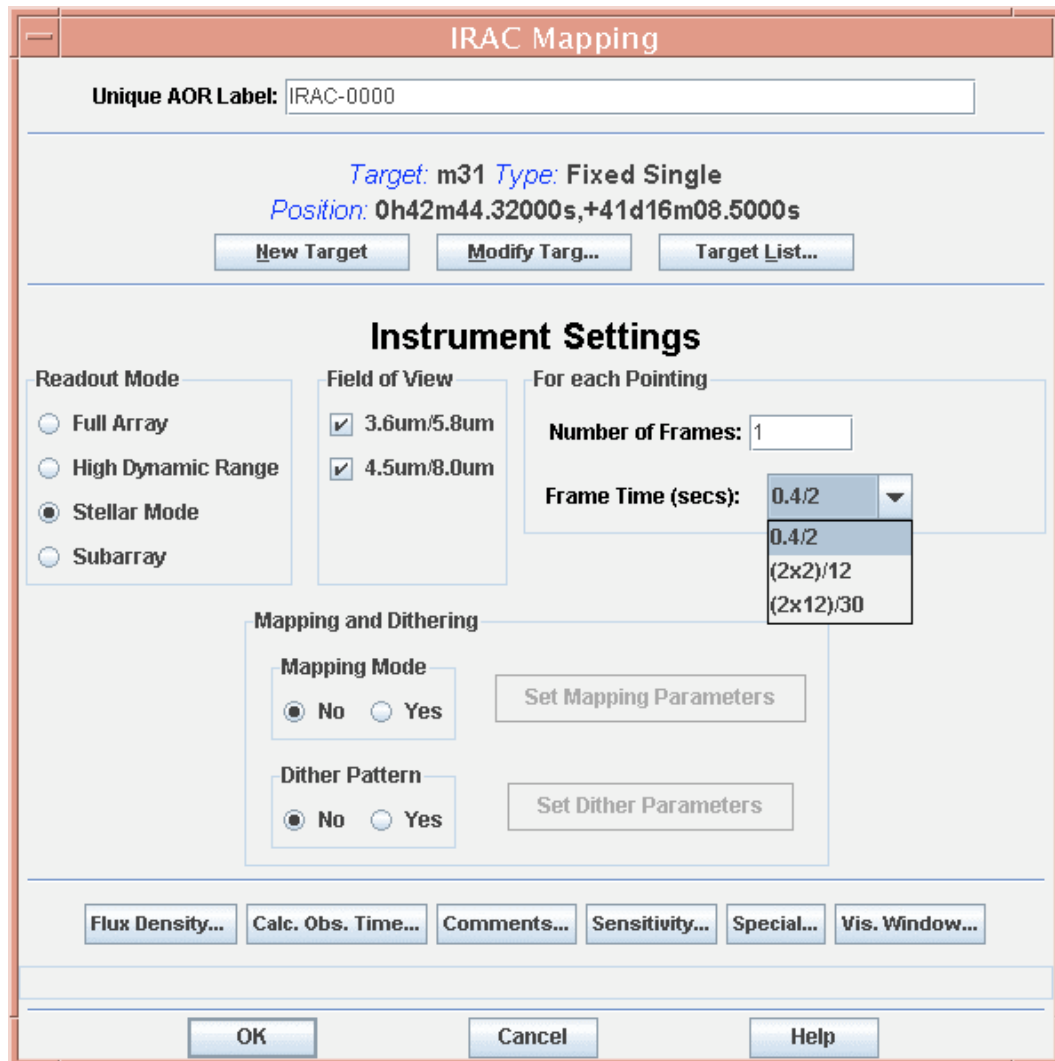
Figure 6.24: IRAC AOT dialog in Spot.

6.2.3.1 Readout Modes and Frame Times

In full-array readout mode, there are 4 selectable frame times: 2, 12, 30, and 100 sec. You should select the frame time that will allow you to reach your target sensitivity while allowing for rejection of cosmic rays, after accumulating all repeat observations of the same piece of sky (by dithers, frame repeats, map

overlaps, and map cycles). The cosmic ray rate is about one hit per array per second.

It is possible that bright sources will saturate the array at the frame time you choose to meet your sensitivity goal. To allow sensitive observations without losing dynamic range, we provide a "high dynamic range" (HDR) option. When this option is selected, the IRAC AOT will take extra frames, with frame times shorter than your selected frame time. The HDR frame times are given in Table 6.9. You can use the shorter frames to measure the brightness of the brighter sources and the longer frames to reach your sensitivity goal. No spacecraft repositioning is done between frames, and the frames always go from shortest to longest. If you elect to perform repeats using the "Number of Frames" entry box on the AOT form, then only the longest frame time is repeated. However, if you elect to perform dithers, then the entire frame set is repeated at each dither position. Independent of the HDR option, taking frame repeats is recommended only in the case of time series measurements. This is because the first frame effect is worse with repeats, and image artifact rejection is best accomplished with dithers. The repeated frames will need additional processing to remove residual bias patterns.

Stellar photometry mode is available for observations of objects much brighter in channels 1 and 2 than in 3 and 4 (typically stars). This mode takes short exposures in channels 1 and 2, and long exposures in channels 3 and 4. Originally developed for engineering observations for taking calibration stars, this mode is available for all observers. Three framesets are available. The shortest set takes a single 0.4 sec frame in channels 1 and 2, and a 2 sec frame in channels 3 and 4. The next set takes two undithered 2 sec frames in channels 1 and 2, and a 12 sec frame in channels 3 and 4. The longest frame time combination available takes two undithered 12 sec frames in channels 1 and 2, and a 30 sec frame in channels 3 and 4. The sensitivities of each frame are identical to those in full array mode. Dithering and mapping are also available in this mode.

For very bright sources, which would saturate the array in full-array readout mode, you can select subarray mode. In this mode, only a small 32×32 pixel portion of the array is read out, so the field of view is only 38″×38″. Further, the telescope must be repositioned for each detector's subarray region to be pointed at the same piece of sky. Mapping is not allowed in subarray mode. However, small maps can be made using a cluster target. Therefore, use subarray mode only for observations of single, very bright targets. Advantages of the subarray mode, in addition to the ability to observe sources that would saturate full-array readouts, are that sets of 64 frames are taken back-to-back, enabling better temporal resolution and decreasing the data volume (which means that rapid sampling can continue for longer periods). Also, because of the more rapid (0.01 sec) sampling, more Fowler samples are taken, resulting in lower readout noise per frame compared to full array mode for the same frame time. In subarray readout mode, there are 3 selectable frame times: 0.02, 0.1, and 0.4 sec. For one commanded image in subarray mode, a set of 64 Fowler-sampled frames are

taken in succession, so that each time an image is commanded in subarray mode, a "cube" of 64×32×32 pixel images is generated. This means that the durations of a single repeat at each of the three subarray frame times are 1.28, 6.4, or 25.6 sec, respectively. The IRAC AOT moves the telescope to point to the subarray region of each requested channel at the target in turn. For the 0.02 sec frame time, data rate limitations allow only data in the channel actually pointing at the target to be taken. For the 0.1 sec and 0.4 sec frame times, data are taken in all four channels at each pointing position, although only one channel at a time will point at the target.

6.2.3.2 Map Grid Definition

If you select "No mapping," then your map grid will consist of a single position at the coordinates specified in the Target section of the AOT. (Dithering can still be applied even if the "No mapping" option is selected.) With "No mapping" selected, if you select both fields of view, first the 4.8/8.0 μm field of view will be pointed at your target, then the telescope will be repositioned so that the 3.6/5.8 μm field of view is pointed at your target. In both cases, data from all 4 arrays are collected, whether they are pointed at your target or not.

Figure 6.25: Map grid dialog box in Spot.

If you select "Yes" under "mapping mode," then you will be presented with a dialog box, shown in Figure 6.25, wherein you specify the rectangular map grid. The map grid may be specified in either "array" or "celestial" coordinates. If you choose array coordinates, then the map grid will be aligned with the edges of the array, such that the map rows and columns will correspond to rows and columns of the array. Specifically, a "column" will be along a line of constant solar elongation, and a "row" will be along an ecliptic parallel [line of constant ecliptic latitude]. It is worth noting that the two IRAC fields of view are at approximately constant solar elongation, so that a map with 1 column and several rows will make a strip along the direction of the separation between the two fields of view and will yield 4-array coverage along part of the strip (if it is long enough). If you select celestial coordinates, then the rows and columns will correspond to J2000 Right Ascension and Declination. You may specify a position angle, degrees E of

N, to orient the raster in Equatorial coordinates. Specifically, if the position angle is zero, a "column" will be along a line of constant Right Ascension, and a "row" will be along a parallel (line of constant Declination). The celestial option can be useful for regions whose shape is clearly defined on the sky. This option also allows two separate AORs, both in celestial coordinates, to have a fully specified location of map grids such that they can be made adjacent to one another independent of observing date (but see below, regarding gaps).

The most important advantage of mapping in array coordinates is that the sky coverage can be custom-tailored, independent of the date when the observation is scheduled. In general, we do not recommend using celestial map grids. A celestial map grid can yield highly non-uniform coverage (including gaps), unless the individual pointings are spaced closely together. There are cases where a celestial grid is preferred, specifically for mapping a highly elongated region. We recommend that all observers considering a celestial grid inspect their survey coverage for a range of possible spacecraft roll angles. Overlaying your AORs on a sky image in Spot can do this. The roll angle itself can be obtained in the Target Entry/Visibility section of Spot, by choosing an observing date when the target is visible. The roll angle is essentially fixed if the source has an absolute ecliptic latitude smaller than 10º (where the total roll angle variation as a function of observing date is less than 8º). Figure 6.26 shows two examples of 3×3 map grids, one in array coordinates and the other in celestial coordinates. For the celestial grid, the spacecraft roll at the time of observation had a position angle of 30º east of north. Not only are there large gaps in the celestial grid, but the coverage is non-uniform in the observed region as well. To eliminate the gaps, you could tighten the spacing of the map grid positions (sacrificing some area covered or some dithering), or you could fix the date of the observation such that the roll angle will be equal to your desired celestial position angle. But be aware that away from the ecliptic poles, not all position angles are available over all positions on the sky.
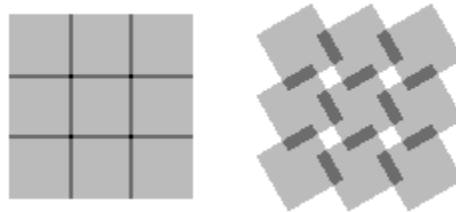
Figure 6.26: Simulated IRAC coverage for a map in array coordinates (left) and celestial coordinates (right).

For large maps, which take more than the 8-hour maximum AOR duration, you will have to break the project into multiple AORs. If your map grid is oriented in array coordinates, then you will not know the boundaries of each AOR precisely, unless you specify the observing date. Therefore, care must be taken in planning and managing large observations. If the large observation can be broken into separate rectangular map grids, they can be specified in individual AORs using

the same central target position, by specifying a "Map center offset" (see Figure 6.25). An example of how to use the map center offsets is given in the Examples section below. In general, to calculate observing times for a large map, generate an AOR that covers a fundamental "tile" for your project, and multiply the observing time of that tile by the number of tiles needed to cover your desired area. Allow for some overlap between tiles, both for calibration and for inefficiencies in the way tiles mesh together.

6.2.3.3 Selecting the IRAC Field of View

As described above (see also Figure 2.1 and Figure 4.5), the four IRAC arrays are fed by two entrance apertures, which are located one above the other in the focal plane with a gap of about 1.5′. For subarray mode, there are four separate fields of view, corresponding to the subarray section of each of the four arrays. The subarray sections of the arrays sharing the same entrance aperture do not cover the same piece of sky, due to an extra reflection. In full-array mode, you select one or both fields of view. In subarray mode, you select one or more of the four subarray fields of view, and the dither pattern will be performed with each of the selected fields of view pointed at the target.
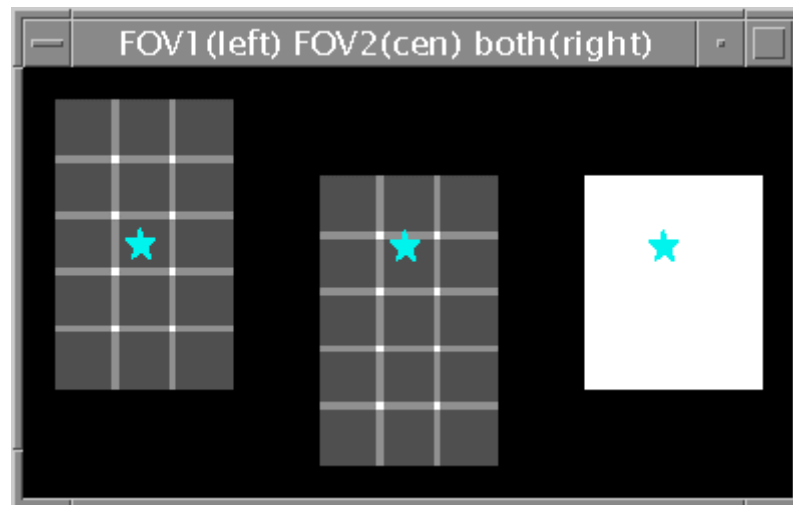
Figure 6.27: Schematic diagram of sky covered by the two IRAC fields of view during a 3×5 map. On the left is the sky covered by the 3.6/5.8 μm field of view; in the middle is the sky covered by the 4.5/8 μm field of view; and on the right is the sky covered by both fields of view. In the left and center panels, lighter shades indicate a higher depth of coverage, while darker shades indicate less coverage. The right-hand panel is white when both fields of view cover the sky.

When making an IRAC map (not allowed in subarray readout mode), you can use the field of view checkboxes to control the way the IRAC AOT performs the map. Whatever your selection, data will be collected with all 4 arrays, and the same sky will always be covered by a pair of arrays sharing an entrance aperture. If you select the 3.6/5.8 μm field of view as primary, then your specified map will be performed with the 3.6/5.8 μm field of view. A nearby portion of sky (significantly overlapping for large maps) will be covered with the other, 4.5/8

μm field of view. This is illustrated in Figure 6.27, where the sky coverage of a 3 column × 5 row map is shown. If you select the 3.6/5.8 field of view, the target position will be centered in the left-hand panel. The serendipitous sky coverage by the 4.5/8 field of view is shown in the center panel, and the sky covered by both fields of view is shown in the right-hand panel.

Selecting both fields of view changes the way the pointing is commanded. If you have selected a "celestial coordinates" map, the AOT performs your map grid, once with the 4.5/8.0 μm field of view centered on each map grid position, then again with the 3.6/5.8 μm field of view centered on each map grid position. This may lead to significant redundancy (and cost in observing time) for large maps. You may prefer to choose a single field of view as your primary field of view, and make your celestial map grid large enough that both fields of view will cover your desired area, regardless of the scheduled date of observation.

If you have selected "array coordinates" and "both" fields of view, then the only special action that will be taken by the AOT is to center the map on the point directly between the two IRAC fields of view. If your map is 1 row by 1 column, the map center will not be imaged by any of the IRAC arrays. Therefore, if you want the same desired area to be covered by all four arrays, you will have to increase the number of rows in the map. For example, suppose that you want to cover an area with a diameter of D, and you want the spacing between map grid positions to be δ. The number of columns in your map is just $1+[(D-A)/δ]$. One IRAC field of view size is $A=5.2'$ and the gap between the two fields of view is $G=1.5'$. The number of rows required to cover this area with both fields of view is therefore $1+(D+G)/δ$. (You should round the number of rows and columns up for any fraction, or else some area at the top and bottom of the desired region will not be covered by all 4 arrays.) For example, if your region has diameter 14′ and you are using a 280″ map grid spacing, then your map should have 3 columns and 5 rows.

There are special cases for which observers should be careful with their choice of field of view. For example, if you want to make a sparse map, with spacing greater than the array size ($δ > A$), then each of the 3 choices of field-of-view (3.6/5.8, 4.5/8, or both) will yield a significantly different map. To make a sparse map with both fields of view, ensuring that all 4 arrays cover the exact same sky, you should split the observation into two AORs, one per field of view. Also, for a very small map ($δ×M∼A$, where M is the number of rows), be sure to make the map large enough that both fields of view will cover the target. If the number of extra rows required using the equations above is excessive, you may prefer to create two AORs, one per field of view.

6.2.3.4 Dither Patterns

Dithering performs a number of functions. It allows the identification and removal of small-scale detector defects such as bad pixels, it reduces noise from the effects of pixel-to-pixel errors in the flat field (in fact for large numbers of

dithers a flat field can be constructed from the data themselves), and sub-pixel dithering can be used to recover some level of information which would otherwise be lost through undersampling of the array. These considerations drive the optimization of the dither patterns in different ways; there is thus no "ideal" dither pattern. For the standard IRAC dither patterns, we have chosen designs that we think should please most observers most of the time. However, the option always remains for observers to define their own dither patterns using array offsets in cluster mode.

For the full-array mode there are two types of dither patterns available. Five such patterns are fixed patterns, which are performed identically at each mapping position. The cycling pattern is a set of dither positions (also referred to as "points"), a different subset of which is performed at each map grid position.

Different patterns are available in subarray mode, as the angular scales covered by the arrays are quite different. Two fixed patterns are available for this mode.

Table 6.14: Characteristics of the dither patterns

Dither Pattern	Scale	Max dither (pixels from (0,0))	Median dither separation (pixels)	Sub-pixel dither pattern
Cycling	Small	11	10.5	½ pixel
	Medium	119	53	½ pixel
	Large	161	97	½ pixel
5-point random	Small	26	23	½ pixel
	Medium	52	46	½ pixel
	Large	105	92	½ pixel
9-point random	Small	16	14	1/3 pixel
	Medium	34	28	1/3 pixel
	Large	69	59	1/3 pixel
12-point Reuleaux	Small	13	15	½ pixel
	Medium	27	30	½ pixel
	Large	55	59	½ pixel
16-point spiral	Small	16	12	¼ pixel
	Medium	32	23	¼ pixel
	Large	64	45	¼ pixel
36-point Reuleaux	Small	17	19	¼ pixel
	Medium	34	39	¼ pixel
	Large	67	78	¼ pixel

6.2.3.4.1   Dithering strategies

The dithering strategy adopted for a particular observation will depend on the type of object and the required depth and resolution of the final image. The characteristics of the available dither patterns are given in Table 6.14. The Reuleaux Triangle patterns were designed with the idea of optimizing the Figure

of Merit of Arendt, Fixsen, & Moseley (2000, Astrophysical Journal, 536, 500). They thus sample a wide range of spatial frequencies in a fairly uniform manner, and are well suited to the Fixsen least-squares flat fielding technique. The 9-point and 16-point patterns were designed to be the optimum size for 1/3 and ¼ sub-pixel dithering, respectively. The "random 9" pattern is based on a uniform random distribution. The "spiral 16" pattern was designed by R. Arendt to provide a pattern which is both compact and has a good figure of merit for self-calibration.

The cycling patterns are designed for AORs having many mapping/dithering observations, but may also be of more general use. The large and medium patterns are Gaussian distributions (with dithers >128 pixels removed). The small pattern is specifically designed for mapping, where only a few dithers are taken at each map position. It is also based on a Gaussian distribution, but the center is downweighted to decrease the fraction of small dithers in the pattern, and it is truncated at a maximum dither of 11 pixels to ensure that maps with up to 280″ spacing have no holes, even if there is only one dither per map point. All the patterns are constrained to have no pair of dithers closer than three pixels in any run of four consecutive points. To use these patterns, you specify the number of dithers to perform at each map grid position and the starting point in the cycling dither table. You may specify a single dither at each map grid position. This will result in a single observation at each map grid position, but the location of that observation will be slightly shifted from the rectangular grid. Such a "fuzzy" grid should allow a better flat field determination and reduce jail-bar or other pattern noise in the resulting mosaic. The cycling dither table will wrap around once the final (311[th]) element is reached. For example, selecting a starting point of 310 and three dithers per map, the first grid position will use points 310, 311 and 1 from the table, the second 2, 3 and 4 and so on. This pattern has a ½ sub-pixel sampling pattern superposed on it, starting with point 1 and repeating continuously every four points (at point 311, the final cycle is simply truncated early, thus patterns which wrap around the table will miss a sub-pixel dither point).

The five-point Gaussian pattern is a general use pattern suitable for shallow observations where the exact sub-pixel sampling is unimportant. It has a ½ sub-pixel pattern, with the 5[th] point at sub-pixel (¼,¼).

Figure 6.28 shows the dither patterns at the default (large) scale. Figure 6.29 shows the cycling dither patterns and the distribution of both the dithers and of the separation between dithers for each scale.

Cosmic ray rejection is greatly facilitated by taking highly-dithered observations. The SSC mosaicking software provides good cosmic ray rejection in cases where each position on the sky is observed with several dithered frames. The SSC mosaicker works well with five or more dithered observations per position on the sky, though will manage some rejection with fewer. It has not been found necessary to take repeated, non-dithered observations at each position in a map or

dither pattern to obtain good cosmic ray rejection if dithering is being carried out with sufficient redundancy.
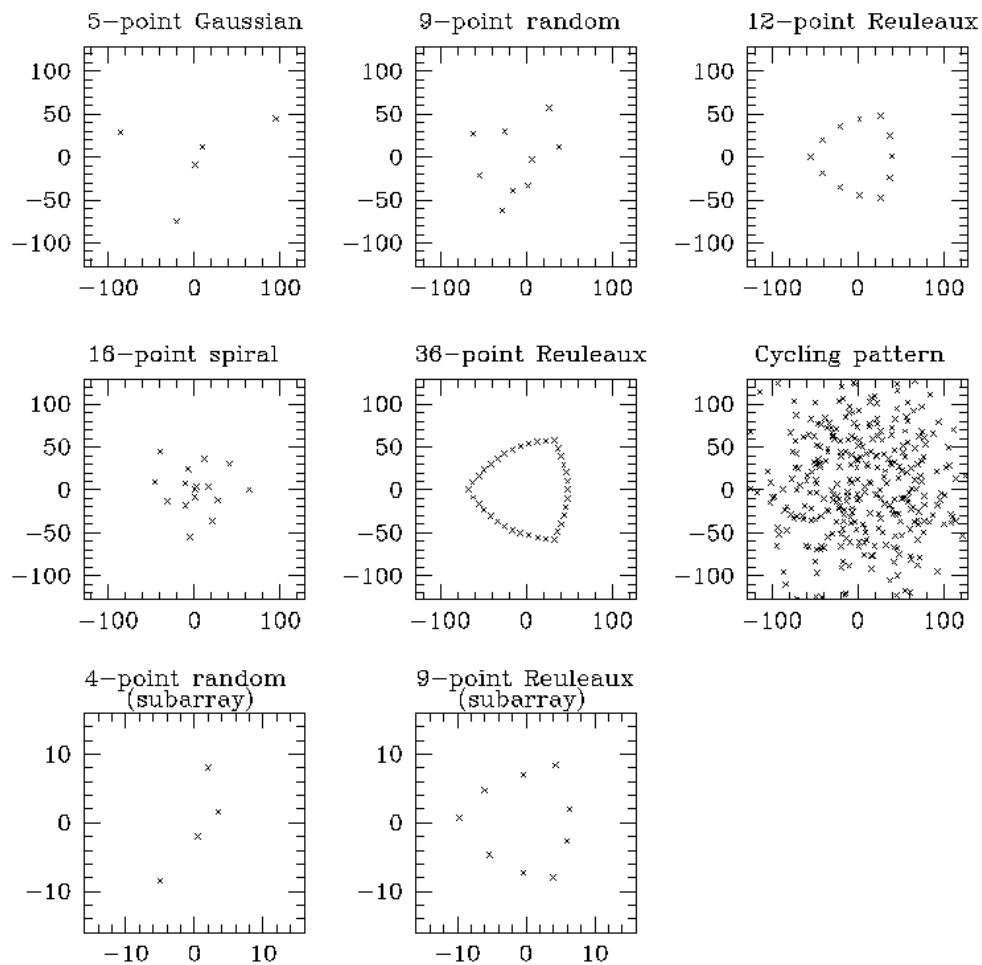
Figure 6.28: IRAC dither patterns for the "large" scale factor.

6.2.3.4.2   Scales and pattern selection in Spot

Each of the IRAC dither patterns is available in three sizes, large (default), medium, and small. For most of the patterns, the scaling of the large, medium, and small patterns is approximately in the ratio 4:2:1. Exceptions are the small cycling pattern, which is about 1/5 of the size of the large cycling pattern and has a lower-weighted inner region to reduce the numbers of small separation dithers, and the 4-point subarray pattern where the scaling is 4:3:1.5. For all the patterns, the sub-pixel dithering is maintained, independent of scale.

The dialog box in Spot for the dither patterns is shown in Figure 6.30, which illustrates how the pattern and scale are selected.
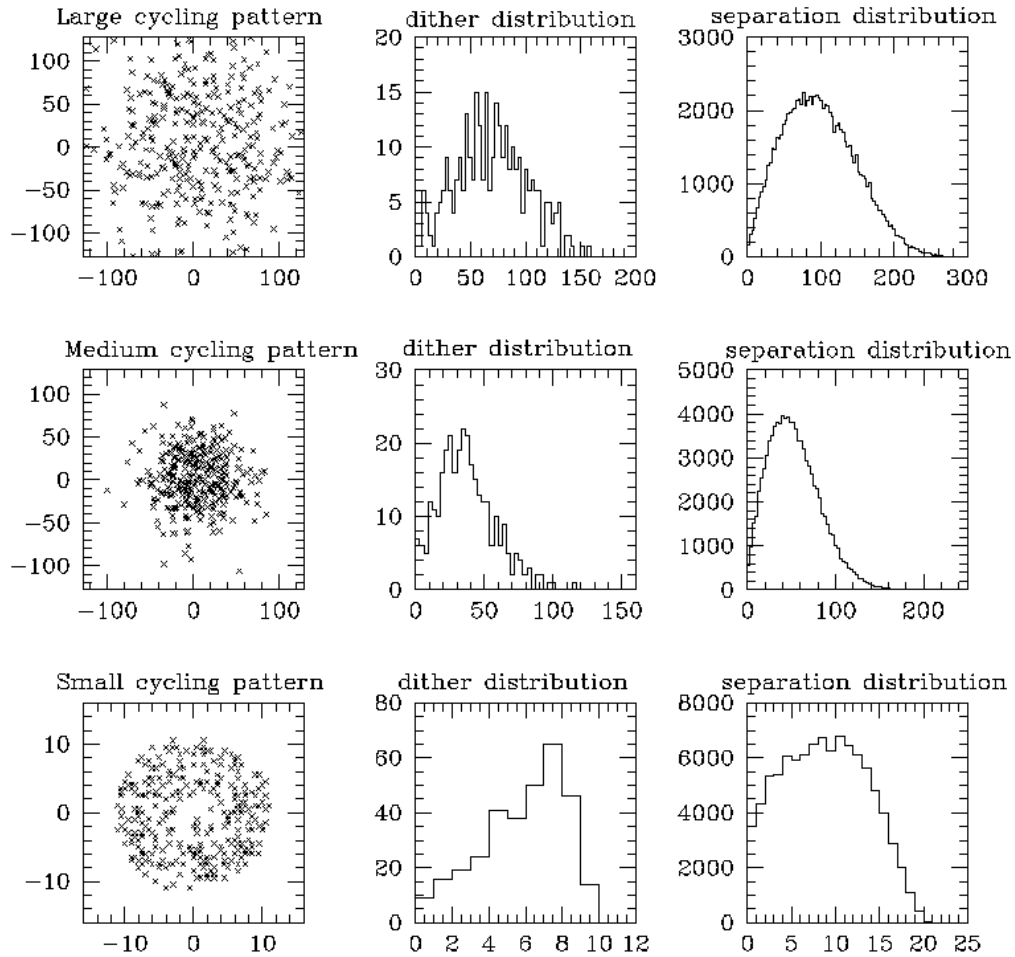
Figure 6.29: Characteristics of the cycling dither patterns, in pixels.
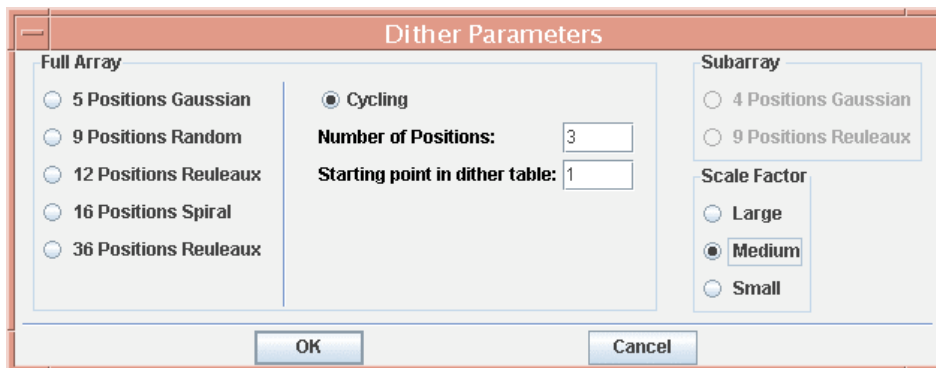
Figure 6.30: Dither pattern dialog box in Spot.

6.2.3.4.3   *Sub-pixel dithering*

Sub-pixel dithering, combined with the drizzle technique (Fruchter & Hook 2002, PASP, 114, 144) to reconstruct the images, can improve the sampling of the mosaics that are obtained from IRAC (or any other) observations. Such strategies

have been used for the WFPC2 and NICMOS instruments on the HST for some time (for details see the HST Drizzle Handbook). Dithering is also needed to calibrate intra-pixel sensitivity variations, and needed for programs requiring accurate photometry and astrometry (Anderson & King 2000, PASP, 112, 1360). To be effective, however, accurate pointing and low image distortion are required. The offsetting accuracy of Spitzer is in the range 0.1″–0.4″. This, combined with the image distortion in the IRAC arrays, places a limit of about ¼ pixel on the sub-sampling that is likely to prove useful in practice. For example, the distortion of the IRAC camera is <1% (see Figure 6.5). Thus for the largest dither patterns, which typically offset up to ±64 pixels from the starting point, the offsets will be up to ±0.6 pixels from the nominal values. Thus only in the small scale patterns, where the offsets are less than ±16 pixels, will the sub-pixel sampling work well, though even on the larger scales some improvement of the images will probably be noticeable.

6.2.4 IRAC AOT Recommendations and Examples

In this section, we present some advice and some completely worked examples of IRAC observations. Further detailed examples are found in the Spitzer Observation Planning Cookbook.

6.2.4.1 Best Observing Practices with IRAC

We recommend the following guidelines when planning IRAC observations. More information on these items can be found below and elsewhere in this Chapter.

1. Always check your planned observation by overlaying it on an appropriate astronomical survey image such as one from the 2MASS/DSS/MSX surveys. Check for the range of allowed observing dates (visibility windows). Verify that bright objects, e.g., from the 2MASS or IRAS catalogs, will not ruin the planned observations in terms of scattered light (section 6.1.2.2.2), persistent images (section 6.1.3.2.2), image artifacts (section 6.1.3.2.10), or saturation (section 6.2.1).
2. Use dithering instead of in-place repeats (see section 6.2.4.1.1 below).
3. Use medium- or large-scale dither patterns unless performing a map with very tight requirements for uniform coverage (see section 6.2.4.1.1 below).
4. If the field you are planning to observe includes very bright sources, consider using the HDR mode to allow for flux measurements of bright sources and identification of saturating sources (section 6.2.4.1.2 below). If you plan to observe a single very bright point source, consider the subarray mode (section 6.2.4.1.4 below).
5. Inspect the field being observed, including the "serendipitous" field covered by one pair of detectors while another is observing the target of interest. Residual images can persist for some time and ruin your observation (section 6.2.4.1.6 below).
6. If rejecting asteroids is important, plan at least two observations of your target, separated by a few hours (section 6.2.4.1.7 below).

7. For the most sensitive observation, use the 100 second frame time. Note that the 200 second frame time is no longer supported; routine sky darks are not taken with 200 second frames.
8. For highest quality images, put your source in the inner 1/3rd of the array.

6.2.4.1.1   Dithering vs. in-place repeats

Dithering is used to eliminate array-dependent or transient artifacts from the true celestial map. A well-dithered map will mitigate the effect of pixel-to-pixel gain differences, which will average down when a celestial-coordinate mosaic is generated. Radiation hit rejection is greatly facilitated by taking highly redundant observations. The SSC mosaicking software provides good radiation hit rejection in cases where each position on the sky is observed with several dithered frames. The medium- and large-scale dither patterns with five or more dithers per position are adequate for this. It is also possible to achieve sufficient redundancy, and accurate maps of large-scale structures, by making a raster map with spacing between rows and columns that is ½ array or less, together with a superposed (small-scale) cycling dither pattern.

Scattered light from bright sources near the edge of the field of view sometimes cannot be avoided, but it can be prevented from contaminating multiple frames by using a dither pattern larger than the characteristic size of the regions that produce stray light. Scattered light is well-rejected by the post-BCD pipeline if dithering and mapping offsets are large enough, for example, with the medium and large dither patterns (but not small), to ensure that stars in subsequent images are moved out of the scattering zones, and the redundancy is high enough to allow for effective outlier rejection (five dithers is sufficient).

We discourage using in-place repeats (successive frames taken at the same position), especially in observations that have less than ten different dither positions, both for the reasons mentioned above and also because the residual first-frame effect will lead to a bias pattern difference between the "repeats" and the "first frames" of a repeat set. Taking in-place repeats is recommended only in the case of time series measurements with stringent requirements for stability, such as observations of planet transits. In general, the better handling of pixel-to-pixel variations using dithers will more than compensate for the reduced amount of integration time; that is, the realized signal-to-noise level will be higher using N-1 dithered observations than N observations with repeats. (In-place repeats are specified in the AOT window under "For Each Pointing/ Number of Frames;" dithers are specified lower in the AOT window under "Dither Pattern.")

Depending on the number of bright sources in your planned field, and the placement of your observation in the schedule relative to other observations that may have bright sources, there may be persistent images in the arrays during your observations. The calibration data (flats and darks) are carefully planned to avoid such artifacts, and they are erased by anneals every 12 hr. If your observation is contaminated by persistent images, the effect on the final data quality is reduced in a dithered observation by at least a factor of 1/N, where N is the number of

dithers. Greater reduction occurs when robust averaging (outlier rejection) is used.

Taking into account possible residuals of the first frame effect and image artifacts due to bright sources, we recommend that sensitive, background-limited maps be made using either as many dithers as possible, or a fine map grid combined with the cycling dither pattern. Shallow, read-noise-limited observations, or raster maps specially designed to identify phenomena with a certain time-dependence (such as asteroid motion), may be best done with combinations of small dither patterns and small map steps. The small-scale dither patterns should only be used in conjunction with mapping using 1/3–2/3 array offsets. The small-scale dither patterns do not provide sufficient redundancy under other circumstances.

6.2.4.1.2   HDR mode

It is possible that bright sources will saturate the array at the frame time you choose to meet your sensitivity goal. To allow sensitive observations without losing dynamic range, we provide a high dynamic range (HDR) option. When this option is selected, IRAC will take extra frames with frame times shorter than your selected frame time. HDR mode is recommended if the target or nearby sources will saturate in the desired frametime. The short frame data are very useful in recovering the photometry of saturated sources, fitting PRFs to saturated sources, and in removing artifacts due to bright saturated sources. HDR mode is only slightly more expensive in observing time than the full frame counterpart.

6.2.4.1.3   Stellar photometry mode

When observing objects that are considerably brighter in channels 1 and 2 than in channels 3 and 4, but not saturating with frame times >0.4 sec, use of the stellar photometry mode is recommended. Short undithered frames are taken in channels 1 and 2 simultaneously with longer frames in channels 3 and 4.

Table 6.15: PRF intensities*

	Radius (″)	2.44	4.88	7.32	9.76	12.20
Ch 1	I/I_0	1.272e-02	6.799e-04	1.390e-04	3.899e-05	1.231e-05
	Enc. Energy	0.855	0.942	0.974	0.991	1.000
Ch 2	I/I_0	1.628e-02	7.040e-04	9.919e-05	6.316e-05	1.928e-05
	Enc. Energy	0.836	0.930	0.968	0.987	0.998
Ch 3	I/I_0	3.236e-02	5.999e-04	2.808e-04	2.936e-05	9.742e-06
	Enc. Energy	0.756	0.931	0.971	0.985	0.991
Ch 4	I/I_0	4.170e-02	1.998e-03	3.224e-04	1.805e-04	2.436e-05
	Enc. Energy	0.634	0.921	0.950	0.979	0.992

* Values are derived from 5× (0.24″/pixel) oversampled PRFs created from observations of calibration stars. I/I_0 is calculated as the flux density/pixel in a 1-pixel annulus at the respective radii, divided by the flux density measured in the central pixel. Enc. Energy is the ratio of the flux density enclosed within the representative apertures to the flux density in a 12.2″–radius aperture.

6.2.4.1.4   Subarray

Subarray mode is useful for observing single, very bright sources, and for obtaining high temporal resolution, since sets of 64 frames are taken back-to-back, with no gap between frames. The saturation limit is much higher in the subarray mode.

6.2.4.1.5   Mapping

In most cases, we recommend mapping in array coordinates since the sky coverage can be custom-tailored, independent of the date when the observation is scheduled. A celestial map grid can yield highly non-uniform coverage (including gaps), unless the individual pointings are spaced closely together. There are cases where a celestial grid is preferred, specifically for mapping a highly elongated region (see the example in section 6.2.4.4). We recommend that all observers considering a celestial grid inspect their survey coverage for a range of possible spacecraft roll angles using Spot visualizations. Be aware that away from the ecliptic poles, not all position angles are available over all positions on the sky. If a large observation can be broken into separate rectangular map grids, as is necessary when the total map duration is more than 8 hours, then individual AORs for each map "tile" can use the same central target position and a "map center offset." Allow for some overlap between tiles, both for calibration and for inefficiencies in the way tiles mesh together.

6.2.4.1.6   Observing near bright targets

If you observe a target next to a very bright target and want to image it with all four arrays, and it will be impossible to avoid having the bright target in one of the arrays (e.g., due to the small roll angle range), it may be better to make the observation in two separate AORs, imaging one channel pair at a time, or to take more frames (with dithers) than would be strictly needed for sensitivity purposes, to beat down persistent images. Remember to justify this in your proposal.

A very common and challenging type of observation is to search for a faint source near a much brighter one. The choice of observing strategy and assessment of the technical feasibility for such a search requires great care. First, it is important to determine the brightness of the point response function (PRF) of the main source at the distance where the putative companion is being sought. Table 6.15 gives the intensity of the PRF, normalized to the peak for a source centered in a pixel, at several distances from the center. Figure 6.31 shows the relative surface brightness of a point source in each channel both along and off a diffraction spike. A well-designed experiment will be able to "cancel" the PRF down to a certain level. For example, if the PRF is down by a factor of one thousand from the peak at the distance being searched, it is plausible to search for features at a level of 1×10^{-4} if some type of PRF removal is done. And, furthermore, having knowledge of the potential shape of the companion (e.g., a point source or disk) will obviously discriminate it from the remaining PRF of the bright source. Remember that the Poisson noise of the bright source at the location of the bright target will always limit the accuracy of the PRF removal.
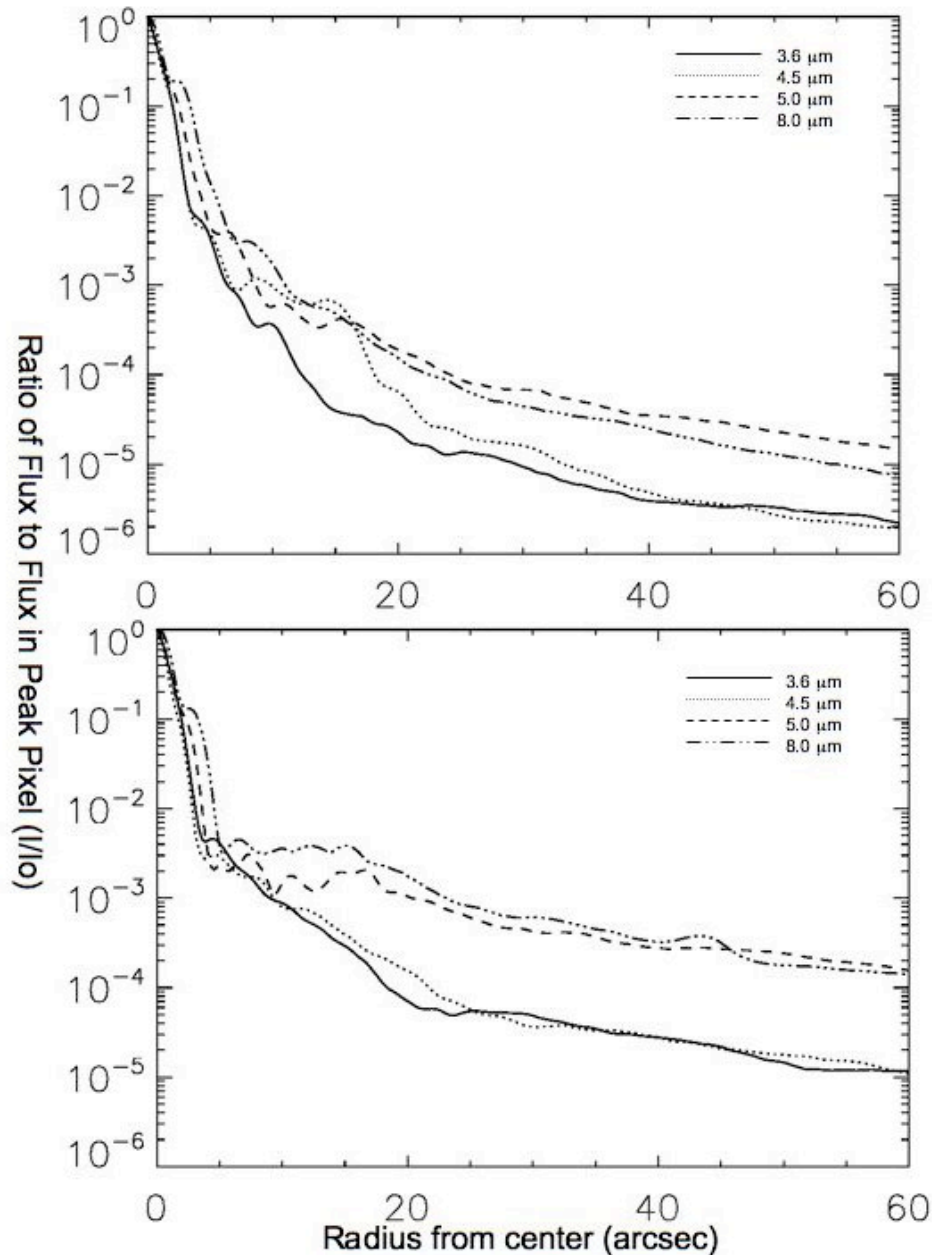
Figure 6.31: Relative surface brightness of a point source for each channel as a function of radius out to 60″ along a diffraction spike (bottom panel) and off a diffraction spike (top panel) as measured with the PRF.

Here is a recommended strategy for this type of observation. Center the bright source on the array, and observe it with a dither pattern with medium scale factor. Use a pattern with good subpixel sampling, to enable future super-resolution techniques; for example, consider the 9-point Random (1/3-pixel sampling) or the 16-point spiral (1/4-pixel sampling). Use the HDR mode, so that there are short frames at the exact same position as the long frames. These may help with centroiding, because they will be less saturated in the core. Use a long enough

frame time for the long frames so that you will gain the required sensitivity on the putative faint source in one (or a few) frames, preferably without saturating the primary source over more than a few pixels. Finally, and probably most importantly, design the experiment as a pair of AORs, with timing constraints such that the source will be seen at two very different roll angles. The range of possible roll angles depends on the ecliptic latitude of the source and must be verified using Spot visualizations near the beginning and end of the visibility windows for a given source. Typically the range of roll angles is about 20 degrees. The PRF will remain fixed in array coordinates, so the basic calibrated images at the two epochs can be subtracted. While any detector artifacts and most of the light from the bright source should subtract out, two images of any real faint object(s) should remain.

6.2.4.1.7   Observing single faint sources

Despite residual image mitigation procedures, observations of a faint source may be affected by long-term residual images from prior observations of bright sources. For observations of single faint sources, we suggest offsetting the target position by $5''-10''$, to reduce the chance of placing a latent from a preceding observation of a bright star that may have used the same dither pattern as your science observation. This offset is most easily accomplished by using a cluster target with array coordinate offsets, entering the offset as the first (and only) cluster position, and then selecting the option to observe offsets only.

6.2.4.1.8   Confusion and other background issues

If you know the accurate (better than $0.6''$) positions of your sources and you are only interested in measuring their fluxes, then the classical confusion limit is not really relevant. It is generally straightforward to align an IRAC image in celestial coordinates to $0.6''$, and often $0.2''$ accuracy can be achieved when there are multiple 2MASS sources in the field.

To separate faint asteroids from more distant targets, it is a good idea to observe your field of view at least twice, separated by a few hours at least, if your source lies at a low ecliptic latitude (within $15°$ of the Ecliptic or so).

6.2.4.1.9   High precision relative photometric monitoring

For observations of exoplanet transits and other sources requiring high precision relative photometry, it is important to minimize systematic effects in the data. In this special case, we recommend observing with repeats and not dithering. By staring for the entire observation, variations in photometry, due to moving the source to other pixels on the array or other positions on a pixel, are minimized. The longest possible frametime that does not saturate the source should be used to maximize the signal compared to the read noise. The field should be inspected to make sure that stray light will not contaminate the observation and that artifacts from bright sources will not overlap the target. If the total duration of the observation is greater than 8 hours, multiple staring AORs should be supplied with close follow-on constraints. A request to transform the multiple observations

into a single instrument engineering request (IER) should be added to the AOR comments and the technical section of your proposal. These observations necessarily produce high data volumes that require the SSC staff to modify the observations so that data is taken only in the desired channel(s).

6.2.4.2 Shallow Survey Example

The goal of a shallow survey is to cover sky rapidly, while maintaining some redundancy in order to reject cosmic rays and reduce effects of pixel-to-pixel gain variations. For this survey, we request coverage of our target region with all 4 IRAC bands. The survey will be conducted with a rectangular grid, with the grid steps aligned with the focal plane array in order to make coverage uniform. We will step by about 95% of the array width for each grid step; specifically we will use a step size of 292.8″ =244 pixels. At each map grid point we will observe 3 dither positions using the cycling dither table. Because we used only a small overlap between map grid positions, we want to constrain the dithers to keep the coverage as uniform as possible, but we also want to separate large-scale photometric variations from sky variations, so we choose the medium dither pattern scale factor as a compromise. The frame time at each position is 30 sec.

We can make an 11×11 map grid, covering 0.7 deg^2, with all 4 arrays, in 4.3 hours as determined using Spot. Except for the small overlap regions between map grid positions and a crust at the edge of the map, the observing time per sky pixel is 3×30 seconds. Using the sensitivity tables above, assuming medium background and a perfect flat field, the 5-σ point source sensitivity is 4.6 µJy at 3.6 µm and 60 µJy at 8 µm.

The observation described here could be used as a "tile" for a survey of a larger area. Suppose you want approximately 9 times the area, so that your survey region could be broken into a 3×3 set of these "tiles." The entire observation could not be done in a single AOR, because it would take longer than the 8 hr maximum duration of an AOR. To implement this large survey, you would generate 9 identical AORs and constrain them to occur within a reasonably short period of time (in order to keep the relative roll angle between the AORs small). Further, you would specify offsets for the center of each AOR. You will need to specify, in addition to the AOR parameters already described, the array coordinate offsets for each AOR that would place them onto the desired grid. Thus if the two array coordinate axes are called (Y, Z) and the desired spacing between map grids is G, then the map center offsets would be (G, G), (G, 0), (G, -G), (0, G), (0, 0), (0, -G), (-G, G), (-G, 0), (-G, -G) for the nine AORs, respectively. Using the constraints editor, you would constrain that the observations all occur within a reasonable time (typically <1 week) of each other, but you should not specify the exact date.

6.2.4.3 Deep Image Example

Suppose we want to make a sensitive image at 8 µm of an object that is less than 4′ in size. In this case, we will not map, but instead will perform many dither

steps in order to minimize the effect of pixel-to-pixel gain variations. For this example, we will use the 36-position Reuleaux triangle dither pattern, medium size, and 100 sec frames. Using Spot we determine that the observation takes 1.1 hours. The 5σ point source sensitivity (medium background) is 9.5 µJy at 8 µm, and the per-pixel surface brightness sensitivity is 0.05 MJy/sr at 8 µm. The same field will also be observed at 4.5 µm, and a neighboring field (not overlapping with the target field) will be observed at 3.6 and 5.8 µm.

6.2.4.4 Elongated Object Example

We wish to make an image of an elongated galaxy oriented 60º E of N. The galaxy has an optical size of 13.5′ by 2.5′, and we want to cover about 17′ with both fields of view. The galaxy is very bright, so we need to use a short frame time (12 sec will work) to avoid saturating or operating exclusively at the high end of the linearity curve. To get the desired sensitivity and source confirmation, we take 5 frames using the small scale cycling dither pattern. In this example, we will perform the observation two different ways, to compare the results.

6.2.4.4.1   Array coordinates

In order to cover the desired area regardless of schedule date, we will need to make the map much larger than the galaxy size. Using the equations above, to cover 17′, we need 5 columns and 4 rows. We select both fields of view to center the image on the nucleus. The duration of the observation is approximately 2328 sec. The map grid size is equal to the array size. The sky coverage is similar to that illustrated in Figure 6.27, with the position angle of the elongated galaxy within the figure depending on the scheduled date. It is possible to make a smaller map, tailored to the size of the target, if we fix the observing date. For a very large observation, fixing the date may be the best solution. But for a small observation such as this one, a better solution is given in the next subsection.

6.2.4.4.2   Celestial coordinates

Another way to observe the elongated galaxy that is both efficient and independent of scheduling constraints is to observe a map grid in celestial coordinates. To get the desired depth of coverage at each position, we could use many different options, but there is one method that will optimize the sky coverage and provide the desired redundancy with minimal overhead. We make a finely spaced map grid that avoids holes due to the (unknown *a priori*) roll angle on the scheduled date and yields the desired 5 observations of each sky position. To do this, we use a map grid spacing of $A/5=61.4″$. The number of rows is 13. The position angle is 60º E of N. We select the cycling dither pattern (small scale) with a depth of one. Making a 13×1 map in celestial coordinates covers the desired long axis of the galaxy, but the perpendicular coverage with the desired depth is only for a relatively narrow strip about the same size as the optical disk.

Therefore, we add 1 extra column to the map and make a 13 by 2 celestial-coordinate map. Figure 6.32 shows the sky coverage by one field of view for this observation. Since both fields of view were selected, the map will be performed

twice, with the galaxy centered first in the 4.5/8.0 field of view and then in the 3.6/5.8 field of view. In Figure 6.32, there are two bands of sky coverage: the upper band is the sky covered in the first map pass by the 4.5/8.0 field of view; the lower band is the "serendipitous" sky covered by the 4.5/8.0 field of view during the second map pass. Using Spot, the duration of the observation is approximately 1500 sec, even with both fields of view selected. This is significantly less than in the array-coordinate map. The savings are entirely due to the celestial coordinate map covering a smaller area, custom-tailored to the desired object size.
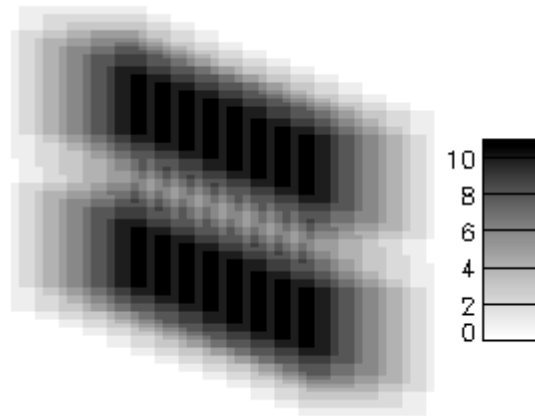
Figure 6.32: Sky coverage for an oversampled, celestial-coordinate map of an edge-on galaxy. The map was made with a position angle 60º E of N. The scale bar gives the number of times each sky position is observed.

6.3 Data

6.3.1 Instrument Calibration

The SSC performs routine calibration of IRAC using observations of standard stars and other astronomical objects. Additional diagnostic data may also be taken at various times to perform certain calibration operations. The data obtained in these observations are used to construct the necessary calibration inputs to the pipeline for the IRAC data processing of science observations. The calibration data files, as well as the pipeline inputs, are available to the general user in the Spitzer archive maintained by the SSC. For more information, consult the IRAC Data Handbook, the Pipeline Description Document, and the paper by Reach et al. (2005, PASP, 117, 978).

6.3.1.1 Astronomical Flux Standards

A number of astronomical standard stars are observed in each instrument campaign to obtain a valid absolute flux calibration. Stars with a range of fluxes are observed at a number of positions across the array and many times throughout the mission to monitor any changes that may occur. Calibration stars with known spectral types and accurate absolutely calibrated fluxes in the IRAC bands have

been determined. These absolute calibration stars are in the continuous viewing zone (CVZ) so that they can be observed at any time necessary and can be monitored throughout the mission.

Four stars are observed in the CVZ at the beginning and end of each instrument campaign. These standards will remain the same throughout the mission, and provide the absolute flux reference for IRAC. Additionally, a (secondary) calibrator near the Ecliptic plane (which will be different for each campaign) will be observed every twelve hours. Its placement in the ecliptic plane is meant to minimize telescope slews. This calibrator is used to monitor any short-term variation in the photometric stability.

Analysis of the flux calibrator data to date indicates that absolute flux calibration is accurate to 3%. Repeatability of measurements of individual stars is good to better than 1.5%, and can be as good as 0.01% with very careful observation design (e.g., Charbonneau et al. 2005, ApJ, 626, 523). The absolute calibration is derived taking several systematic effects into account. The steps are described in detail by Reach et al. (2005, PASP, 117, 978). If this methodology is not applied, then point source photometry from the BCDs can be in error by up to 10%.

6.3.1.2 Sky Flats

To get the most accurate measure of the full system gain, including the effects of the telescope and the IRAC pickoff mirrors, one must use observations of the sky.

This is done using many dithered observations of a network of 22 high zodiacal background regions of the sky in the ecliptic plane, which ensures a relatively uniform illumination with a reasonable amount of flux. One such region is observed in each campaign. The data are combined with object identification and outlier rejection, producing a product analogous to a "median sky flat," as is commonly constructed during ground-based observing, and which will be an image of the presumably smooth celestial background, further smoothed by the dither pattern. The resulting flat field will be divided into the science data. Pixel-to-pixel accuracy of the flat-fielding derived from a single observing campaign is 2.4%, 1.2%, 1.0%, and 0.3%, 1σ, for bands 1 through 4, respectively. Using combined flats ("super sky flat") from the first two years, the estimated pixel-to-pixel accuracy is 0.5, 0.2, 0.2, and 0.05% in channels 1-4, respectively. Users should note that the flat field data are generated from the zodiacal background, and are appropriate for objects with that color. There is a significant color term, of order 5%-10%, for objects with a Rayleigh-Jeans spectrum in the mid-infrared (such as stars); see the IRAC Data Handbook for more information.

Note that for deep survey observations and other data sets with a large number of frames and a good dithering strategy, the system gain could be determined by the actual survey frames themselves, rather than using the standard set of dedicated observations of some other part of the sky. The dither and mapping pattern need to be specified to optimally relate each pixel in the array to all the others, as well as mapping out the region at the required sensitivity. Tests of self-calibration

with suitable datasets have shown that the SSC pipeline results are comparable. The SSC will not, as a matter of course, undertake such special processing (self-calibration) as part of the automated pipeline. The SSC-generated data products will always use the dedicated calibration data.

6.3.1.3 Skydarks

Dark current and bias offsets are calibrated via the standard ground-based technique of dark subtraction. As part of routine operations, the SSC observes a dark region of the sky ("skydark") near the north ecliptic pole at least twice per campaign (at the beginning and end). These data are reduced and combined in such a way as to reject stars and other astronomical objects with size-scales smaller than the IRAC array. The resulting image of the minimal uniform sky background contains both the bias and dark current. When subtracted from the routine science data, this eliminates both of these instrumental signatures. Naturally, this also subtracts a component of the true celestial background. The SSC includes a COBE-based model estimate of the true celestial background, which is the same as that returned by Spot. *Note that our lack of an isolated measurement of the dark current and bias offset during shutterless operations limits the ability of IRAC to measure the true celestial background.*

6.3.1.4 Distortion and PSF Map

The PSF over the field of view of each of the four bands has been characterized, and the optical distortion has been measured. This has been done by observing an open star cluster for which we have good ground-based astrometry. The distortion is included in the WCS header keywords. The positions of both IRAC fields of view have also been determined relative to the PCRS and other science instruments.

6.3.1.5 Linearization

Non-linearities in the IRAC detectors were extensively calibrated during ground-testing using special ground support equipment and also in flight using extended astronomical objects. These calibration data allow the detectors to be linearized to 1% for 90% of their full-well capacity. The linearity solutions are spot-checked during flight, on timescales of many months, via observations of bright, extended objects such as elliptical galaxies. If we find a significant deviation from the previous linearity characterization, new solutions will be derived.

6.3.1.6 Routine (Frequent) Calibrations for Normal Observations

Some calibrations occur extremely infrequently. Examples of these are linearization, distortion mapping, etc. Other calibration measurements are repeated more frequently, and in particular at least every time IRAC is powered on. These calibration observations will be performed at specific intervals as required, rather than having specific calibration observations for each individual data set. Every 12–24 hours, observations of secondary calibration stars will be

obtained, for example. Skyflats and darks are derived at the beginning and end of each campaign.

6.3.2 Data Products

This section describes the basic data products the observer will receive from the SSC. More details are contained in (a) the *IRAC Pipeline Description Document,* which describes the algorithms implemented in the IRAC pipeline and is available on the SSC website; and (b) the *IRAC Data Handbook,* which provides a detailed guide to the data products and how to use them, and is available on the SSC website. The data delivery consists of a directory hierarchy with a name unique to that AOR. In this hierarchy are the BCD data, as well as a number of subdirectories containing the raw data, the calibration files, log files, and the post-BCD data. The exact contents of the data delivery vary according to what the observer has requested from the SSC data archive via Leopard. IRAC data are supplied as standard FITS files.

Each file consists of a single data collection event (i.e., a single exposure), and contains one image corresponding to one of the four IRAC arrays (the exception being post-BCD products, described below). The FITS headers are populated with keywords including (but not limited to) physical sky coordinates and dimensions, a photometric solution, details of the instrument and spacecraft including telemetry when the data were taken, and the steps taken during pipeline processing. There are three primary image data types that are supplied for each AOR (raw, BCD, mosaic).

6.3.2.1 Raw Data

Raw data are wholly unprocessed except for those steps necessary to render them into a readable FITS format, i.e., depacketization and decompression. This is the form in which data enters the IRAC pipeline. Most observers are unlikely to use these data beyond sanity checking of the pipeline data products. They are, however, supplied in the event that observers wish to reprocess their IRAC data in a different manner from the SSC pipeline. By comparison to ground-based astronomy, these are the raw data one gets from a camera and writes to disk while observing at the telescope. The raw data consist of 16-bit integer FITS files. The headers are populated with all of the ancillary science telemetry keywords.

6.3.2.2 Basic Calibrated Data, or BCD

BCDs are exposure-level data after having passed through the IRAC pipeline. There is one BCD for each integration taken by IRAC. Instrumental signatures have been removed (as described in section 6.3.3.2), and the BCDs are absolutely calibrated into physical units (i.e., MJy/sr = 10^{-17} erg s^{-1} cm^{-2} Hz^{-1} sr^{-1}). In addition, the pointing refinement (see section 6.3.2.4) based upon 2MASS catalogs will be supplied in the BCD header. Continuing the analogy with ground-based observing, the BCDs are data that have been reduced, but not yet combined into a final image. This is the primary science data product produced by the SSC.

Ancillary files are supplied with each BCD (see Table 6.16). These ancillary files contain several types of information regarding each pixel in each image. An image containing the uncertainty for each pixel is supplied. A mask image contains status bits indicating the probability that any given pixel has been affected by a radiation hit; whether or not the linearity solution could be applied; if a pixel is saturated, dead, hot (always on), or abnormally noisy. Pixels with strong residual images are flagged in the mask file. Log files are also supplied, and from these and the header keywords the entire pedigree of every data product can be derived.

Table 6.16: IRAC Basic Calibrated Data (BCD) pipeline products

Data	Contents
BCD	Basic Calibrated Data - dark-subtracted, linearized, flat fielded, and flux-calibrated among other things (see section 6.3.3.2)
Uncertainty Image	Array of calculated uncertainties for the BCD
Imask	Data quality flag image for the given science DCE
Residual Image pixels	Bit that is set to indicate residual image prediction
Radiation hit pixels	Bit that is set to indicate the probability of radiation hit event for a given pixel

All of the data and ancillary files are in FITS format, containing a header with keywords and their values followed by a binary image, except certain log files that are in simple ASCII format. The standard FITS header keywords are all present, so that essentially any FITS file reader is able to read the files, and all images have 256×256 pixels. Subarray data are in a cube (32×32×64 pixels). Observers should expect to receive a data volume of approximately 3.5 Mbytes per frame as their BCD, including ancillary files as described in Table 6.16. Information about the data that observers receive is posted on the SSC website under "Data Archives/Analysis" or in the "Data" section under "IRAC."

6.3.2.3 Calibration Files

For each BCD, the pipeline calibration server generates several estimates of the current detector characteristics. These include a map of the pixel-to-pixel response variation (flat field) and an estimate of the dark bias consisting of a laboratory dark and a sky dark. These calibration files are supplied to the observer. Observers are also able to request from the Archive the files that were used to generate the sky darks, sky flats, and absolute calibration. The photometric calibrators are not included with each science observation, but they are available via a separate request to the Archive.

6.3.2.4 Extended Pipeline Products (Post-BCD Pipeline)

Pipeline processing of IRAC data by the SSC also includes more advanced processing of many individual IRAC frames together to form more "reduced" data products. Known by the generic title of "post-BCD" processing, this extended pipeline refines the telescope pointing, attempts to correct for residual

bias variations and produces mosaicked images. We do not attempt to improve (relative to the BCD) the point source or extended emission flux calibration by automatically comparing to a reference source catalog. The mosaic only includes data from a single AOR, so that observers who break their map into multiple AORs (for example those whose maps cannot be completed within the 8 hour time limit of an IRAC AOR, or those who are making multiple-epoch observations) need to recombine them to obtain their ultimate images.

All IRAC BCD images contain a pointing estimate based on the output of the Spitzer pointing control system (startracker and gyros), i.e., the boresight pointing history file. This initial pointing estimate is accurate to about $0.5''$. The post-BCD pipeline performs additional pointing refinement for all IRAC frames. This is achieved by running the SSC point source detector on the channel 1 and 2 frames and comparing the resultant list of point sources to the 2MASS catalog. The results are then averaged, and the known focal plane offsets between all four channels are applied to produce a "superboresight" pointing history file, which is then applied to the data during end-of-campaign reprocessing. This improves the pointing accuracy of the frame to better than about $0.3''$. This refined RA, Dec appears in the header as the CRVAL1, CRVAL2 keyword values.

The pipeline SSC mosaicker produces a single image (1 per band) from many input images. First, the BCDs are corrected for overlap consistency. The parts of the images that overlap are forced to have the same background value via addition of an offset. Then a "fiducial frame" is derived. This is the definition for the output frame in terms of its physical size, projection, and orientation. Because IRAC has such a large field of view, projection effects are non-negligible, and the mosaicking and coadding process must reproject the data. The fiducial frame finder seeks to minimize the amount of "blank" area in the output mosaic by rotating the output projection such that it is aligned with the map axes. This is useful for long thin maps, where potentially the output mosaic could be very large, but with a great deal of empty space. The mosaicker then reprojects all of the input data onto the output projection. It reads the SSC WCS, which contains the field pointing center, rotation, scale, and instrument distortion, and reprojects this onto a standard TAN FITS projection. In the process, the data are undistorted. The reprojected images are interpolated onto the fiducial image frame with outlier rejection, rejecting radiation hits that happen in overlapping observations. The outlier rejection scheme is specifically designed to work well in the case of intermediate coverage and may not be adequate for all observations and science programs. In addition to a sky map (in units of surface brightness), a noise image and coverage map are also produced.

The post-BCD pipeline modules have been made available for general public use by the SSC. They consist of a number of C-modules connected via PERL wrapper scripts. Namelists are used for input. In most cases their operation simply consists of supplying the software with a list of input image files; by default they read and understand the IRAC image headers.

6.3.3 Data Processing

6.3.3.1 Overview

The main processing of the IRAC science data takes place in the "BCD pipeline," which is a set of modules that work in tandem with calibration and pointing servers. The BCD pipeline takes as input the raw data files, and generates the Basic Calibrated Data (BCD) as its final output. The pipeline is run automatically upon receipt of individual raw files. The BCD pipeline is triggered for all science frames generated as part of an AOR. Calibration data are sent to other pipelines for processing, and their output is made available for the calibration server. When the BCD pipeline is running and requires a calibration file (such as a dark current image), a query is sent to the calibration server, which returns the appropriate file. The calibrated images (BCDs) from an AOR are combined in the post-BCD pipeline to create a mosaic. The extended pipeline is exported for observers, so that they can combine data from multiple AORs. The BCD pipeline is not exported, because of its inextricable link to the Spitzer Science Operations DataBase (SODB).

6.3.3.2 Science DCE Processing

The following is an explanation of the steps taken by the IRAC BCD pipeline (Figure 6.33). Further guidance to IRAC data reduction is provided in the IRAC Data Handbook and the IRAC Pipeline Description Document.

1. The header of the raw data file is validated to ensure that it is an uncorrupted IRAC image with the shutter open and lamps off.
2. Several computed keywords (such as exposure time, commanded instrument and telescope parameters, and the pointing information) are added to the image from the telemetry and the SODB. Many of the detailed IRAC telemetry channels are removed from the header; they are still present in the raw header.
3. The 16-bit integer FITS images are received and then rectified so that increasing flux produces increasing DN (which is not the case initially for the IRAC InSb arrays).
4. The 16-bit integer FITS images are converted to 32-bit floating-point FITS images.
5. Detector wraparound is detected and corrected. IRAC can generate negative flux numbers as a result of Fowler sampling. Although IRAC uses a signed 24-bit internal data representation, it only sends unsigned 16-bit integer data and in the process of doing so discards the sign bit. The wrap detection module uses the knowledge that very high DN values exceed the full-well capacity of the detector, and hence they are actually wrapped negatives. Fowler number, barrel-shift, and gain have been chosen such that the entire full-well capacity of the detectors can be represented by 16 bits.

6. The raw data units are renormalized to account for IRAC Fowler sampling and bit-shifting. Observers need not take Fowler sampling into account when reducing their data.
7. Due to the design of the IRAC electronics, there is a maximum rate at which they can react to changes in pixel intensity during readout. Thus, it is not possible to go from full to zero pixel intensity between two adjacent pixels in a single readout channel (equivalent to "pixel latency"). This artifact appears as a decaying ghost following a bright pixel, along the fast readout channel, which is horizontal in the IRAC images. This bandwidth limitation is corrected, based on measurements of the detector bandwidth in each readout channel taken during IOC/SV.
8. The dark current and bias offset are subtracted using a laboratory dark, which should be a good first-order approximation of the dark current.
9. First-frame effect (timing-dependent bias offset) is subtracted. These offsets have been extensively calibrated on the ground as a function of array timing. These offsets are a constant level per readout channel, and thus they can be subtracted from the images.
10. Multiplexer-bleeding correction is applied. The InSb arrays exhibit an effect similar to the electronic bandwidth artifact known as "muxbleed." However, it has a different characteristic curve and is corrected via a series of look-up tables. The correction is not complete, and post-BCD processing is required to reduce the artifacts.
11. Data are corrected for non-linearity. A function (quadratic) was fit individually to the response function of every pixel based on laboratory and flight calibration. This function is then used to linearize the response of the individual pixels. Saturated pixels are identified and marked in a data quality file. The processed pixels are linear to better than 1%.
12. Scattered light patterns are removed from data. A "butterfly" pattern (for channels 1 and 2) and a "tic-tac-toe" pattern (for channels 3 and 4) caused by internal scattered light are scaled using the estimated zodiacal background level and subtracted from each image. They are subtracted from all calibration data as well.
13. The "skydark" is subtracted from the image; this is a correction to the bias relative to the laboratory dark, which is monitored throughout the mission using dark sky patches. It is chosen from the nearest calibration in time.
14. Radiation hits (cosmic rays) are detected at the single frame level using a median filter. Radiation hits are not removed from the BCD. Instead, a probability that a given pixel is affected by a radiation hit is noted in the data quality file (imask). The primary method for removing radiation hits is in the post-BCD pipeline, which uses robust averaging of observations of the same piece of sky.
15. The data are flat fielded. The flat field image is generated from observations of bright sky patches.
16. The channel 1 and 2 images are flipped so that they have the same orientation as channels 3 and 4. This transposition is a result of the

dichroic beamsplitters. As a part of this operation, associated mask files are flipped as well.

17. The photometric calibration, appropriate for point sources, is applied. The pixels are in units of surface brightness (MJy/sr).

18. All of the DCEs from the present AOR are used to determine where there are potential residual images. The data pipeline does not correct the BCD for residual images, since this correction may be quite complex and dependent on a number of different factors, and hence cannot be easily executed in an automated fashion. However, we determine which pixels are likely to produce residual images above the local noise level, and we flag those pixels in subsequent DCEs based on the residual image decay curve.

19. Finally, the 2MASS data are used to refine the pointing. The refined pointing keywords are added to the header of the final BCD product at this stage.

In the future, we plan to implement steps in the pipeline to correct banding (a cross stripe centered on a bright pixel in channels 3 and 4) and column pulldown (bias depression in columns containing bright pixels in channels 1 and 2).

At this stage, the BCD has been produced and consists of a single image corresponding to each exposure made by IRAC. The header of the image contains the standard FITS keywords and the image distortion coefficients, as well as sufficient information to trace the calibration and pipeline pedigree. This image is calibrated in physical units and has a pointing solution attached.

Further processing to generate the extended pipeline products mentioned above will use the BCD files as input. Post-BCD pipeline modules are available on the SSC website to match the background brightnesses of overlapping images, take care of column pulldown (see the IRAC Data Handbook for more information), correct for muxbleed, correct for banding, register images (relative to each other) using sources in the overlapping region, and generate mosaics from all the images in an AOR. Another set of pipeline modules is under development to extract point sources from the images and measure their fluxes (using dedicated observations of the point spread function), place the fluxes into source lists for each wavelength, and merge the source lists into a catalog. The algorithms have been designed to be robust and strive for completeness at the 10σ level, thus they do not reach the ultimate sensitivity of the data.

6.3.3.3 Data Reduction Software

6.3.3.3.1   Data Reduction

The actual BCD pipeline will not be released, due in large part to the degree to which it is tied to SSC-specific infrastructure. However, detailed documentation

on the algorithms used by these modules is in the IRAC Pipeline Description Document.

The software modules that make the post-BCD products (mosaicker/coadder, point source extractor, and bandmerger) have been released to the general community and are available from the SSC website. This software consists of modules written in C with PERL wrapper scripts to control them. The primary advantage they offer over existing astronomical software packages is the ability to read the IRAC ancillary data (mask files, etc.) during processing, and a robust handling of image distortion and reprojection. The latter is important both due to the non-negligible distortion in the IRAC detectors and their large areal coverage. The SSC has encoded image distortion using the TAN-SIP projection type, which is now supported by several software packages. It is anticipated that the community will also find this software useful in processing multiple related AORs, which the SSC does not do automatically.

6.3.3.3.2   *Data Analysis*

Most scientific analysis, such as optimal mosaicking and source extraction, require additional processing beyond the Basic Calibrated Data. Since data products are delivered in standard FITS format, observers may use any of a large variety of data processing packages available to the general community (e.g., IRAF, IDL, etc). The SSC pipeline provides mosaicked images of entire AORs, and the part of the pipeline that takes the BCDs and combines them into the mosaic is provided to observers. Users should be able to use the post-BCD software themselves for additional, multi-AOR data reduction, or to improve the quality of mosaics as necessary for the desired science goals. The software is designed without arbitrary limits to the size of datasets that it can process.
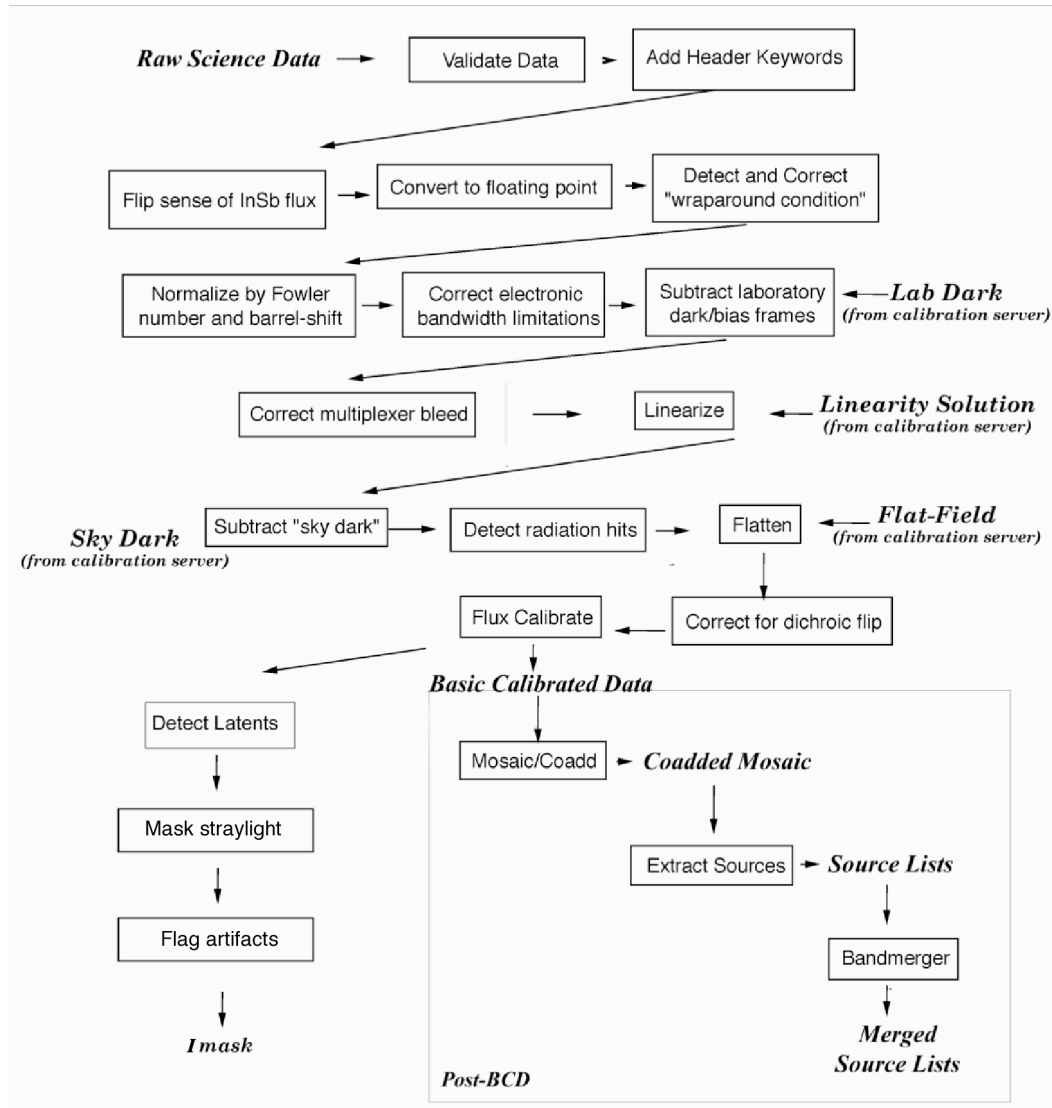
Figure 6.33: IRAC pipeline.

7 Infrared Spectrograph (IRS)

7.1 Instrument Description

7.1.1 Overview – What is the IRS?

The Infrared Spectrograph (IRS) provides the Spitzer Space Telescope with low and moderate resolution spectroscopic capabilities from 5.2 to 38.0 μm. The IRS is composed of four separate modules, with two modules providing R ~ 60–120 spectral resolution over 5.2–38.0 μm and two modules providing R ~ 600 spectral resolution over 9.9–37.2 μm (see Table 7.1). The names of these modules are summarized (along with frequent abbreviations) in Table 7.2. The IRS has no moving parts and each module has its own entrance slit in the focal plane; see Figure 2.1 and Figure 4.5. The low-resolution modules employ long slit designs that allow both spectral and one-dimensional spatial information to be acquired simultaneously on the same detector array. Two small imaging sub-arrays ("peak-up arrays") in the Short-Low (SL) module allow objects to be placed accurately into any of the IRS entrance slits and provide images of sources for photometric or structural information. The high-resolution modules use a cross-dispersed echelle design that allows broad spectral coverage in a single exposure. The median 1σ continuum sensitivity for the IRS low-resolution modules is about 0.06 mJy from 6 to 15 μm, and 0.4 mJy from 14 to 38 μm in 512 seconds of integration with low background. The median 1σ line sensitivity estimates for the short- and long-wavelength high-resolution modules are about 2.7×10^{-19} W m^{-2} and 8.5×10^{-19} W m^{-2}, respectively, for 512 seconds of integration with low background (see Figure 7.18 through Figure 7.22). The IRS was designed to achieve high sensitivity at the expense of reduced dynamic range. Therefore, observers should take great care in interpreting weak spectral features.

The major subsystems of the IRS are:

- The *Optical/Mechanical* subsystem consisting of the optical elements and their supports.
- The *Stimulation* subsystem consisting of the flood illuminators and controls. These are used for in-orbit calibration.
- The *Thermal* subsystem consisting of the temperature-sensing elements, heater, heat control circuits, and heat straps to the cryogenic attachments used to control the array temperature.
- The *Detector* subsystem consisting of the focal plane arrays and their associated multiplexers, cabling, and warm preamplifiers/signal conditioning circuits.
- The *Electronic* subsystem consisting of the warm control electronics (which also support MIPS), power conditioning circuits, and the interface electronics that receive power and commands from the Observatory and pass detector outputs to the Observatory for transmission to Earth. This subsystem includes the Instrument Computer, which converts commands into drive voltages and clocking tables, loads them into the warm control

electronics, collects digitized data, performs simple arithmetic on frame stacks and computes telescope offsets from peak-up images when necessary.

Table 7.1: IRS module operating wavelengths

Module	Detector	Wavelength Range (μm)	Resolving Power
Short-Low (SL)	Si:As	5.2–8.7 (SL 2nd, bonus)	60-127*
		7.4–14.5 (SL 1st)	61-120*
Blue Peak-Up (SL)	Si:As	(13.5–18.7)	(~3)
Red Peak-Up (SL)	Si:As	(18.5–26.0)	(~3)
Long-Low (LL)	Si:Sb	14.0–21.3 (LL 2nd, bonus)	57-126*
		19.5–38.0 (LL 1st)	58-112*
Short-High (SH)	Si:As	9.9–19.6	≈600
Long-High (LH)	Si:Sb	18.7–37.2	≈600

* $\Delta\lambda$ is approximately constant as a function of λ for each module and order.

Table 7.2: IRS slit description synonyms

Formal name	Abbreviation
Hi Res 9.9-19.6 μm	Short High (or SH)
Hi Res 18.7-37.2 μm	Long High (or LH)
Low Res 5.2-8.7 μm	Short Low 2nd order (or SL 2nd, or SL2)
Low Res 7.4-14.5 μm	Short Low 1st order (or SL 1st, or SL1)
Low Res 14.0-21.3 μm	Long Low 2nd order (or LL 2nd, or LL2)
Low Res 19.5-38.0 μm	Long Low 1st order (or LL 1st, or LL1)

> **A brief summary of the IRS appears in the First Observations with the Spitzer Space Telescope special issue of the Astrophysical Journal Supplement Series (see Houck et al. 2004, ApJS, 154, 18). A copy of this paper is available on the SSC website.**

7.1.2 The IRS Cold Assembly

The IRS Cold Assembly (i.e., the four spectrograph modules, the cold interface board, the IRS baseplate, and the associated cables) is mounted directly onto the Spitzer facility base plate. The instrument itself is comprised of four spectrograph modules whose packaging concept is one of "bolt and go" – the modules have no adjustments other than shimming of the focal plane mount assembly in order to achieve proper focus.

The IRS interfaces to a heat sink temperature of 1.4 K–1.6 K. During operation, the Focal Plane Array operating temperature is tightly controlled at temperatures a few degrees above the interface temperature. Once the observing run has begun, the temperature stability of the array is controlled over the entire observing cycle. Specific details concerning array temperature control are defined in section

7.1.4.2 of this document.

Each IRS spectrograph module has its own enclosure and internal baffling and is attached to the common IRS baseplate at three points. The common baseplate is in turn secured by a three-point attachment to the floor of the Spitzer Multiple Instrument Chamber (MIC). In order to provide sufficient shielding to meet the ionizing radiation requirements of the IRS detectors, there is a minimum path length of 1.8 cm through aluminum, along any direction into each focal plane. Figure 7.1 shows a diagram of the four modules installed on the baseplate along with the cold interface board (see section 7.1.4).
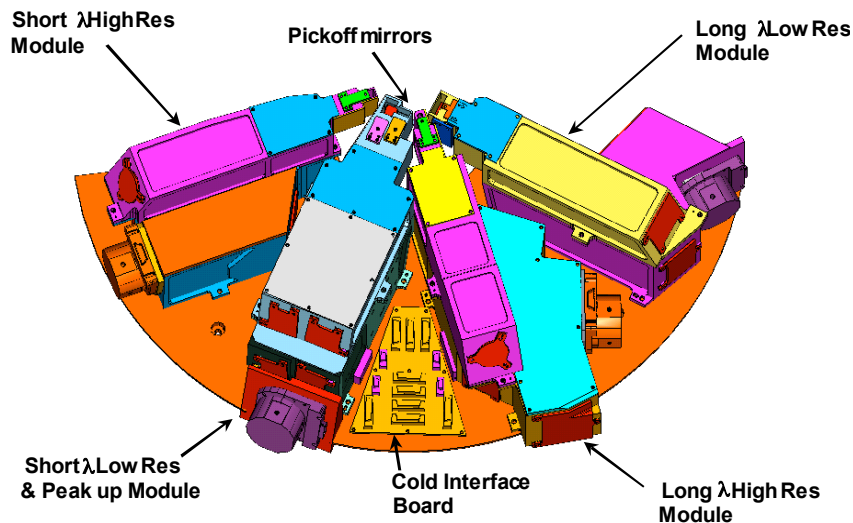
Figure 7.1: The IRS Cold Assembly.

7.1.3 Combined Electronics

7.1.3.1 Hardware

The architecture of the Combined Electronics (CE) subsystem incorporates two independent and redundant electronics boxes (A & B). The CE subsystem is designed so that it will meet all of its performance requirements within 30 minutes of being turned on. Due to the different locations for the A and B sides in the spacecraft bus, side A can meet its performance requirements more rapidly after turn on than side B. This, together with other operational advantages in using side A, indicate that this side will be baselined for use, with side B only being used in case of a fault in the A side.

The CE subsystem is used to operate both the IRS and MIPS cold instruments, although only one instrument can be used, or even powered up, at a time. While using the IRS, separate circuit boards in the CE handle detector array clocking, detector array signal conditioning, communication with the spacecraft, command and data processing (using a RAD 6000 CPU), and engineering data collection functions. When the IRS is operating, all four detector arrays are clocked simultaneously. The on-chip output drivers are independently powered, so that

any number of the IRS arrays can be powered on simultaneously. However, the array outputs are multiplexed to the analog signal conditioning circuit so that it is only possible to capture data from one array at a time.

7.1.3.1.1 Analog Signal Conditioning

The Si Analog Signal Processing electronics select and condition the signals from the science arrays in use. Figure 7.2 shows a block diagram of the signal path for the output from the Si detector(s). The Si Analog Signal Processing electronics provide signal output and processing for up to five (4 IRS plus 1 MIPS) Si arrays on a 5×4 channel multiplexer with four channels per array. The output signal is 1.2 μV per electron; digital saturation occurs at 301,466 electrons. Detector saturation occurs at much higher values. The voltage becomes more negative with increasing detector charge (see Figure 7.3 and Figure 7.4).
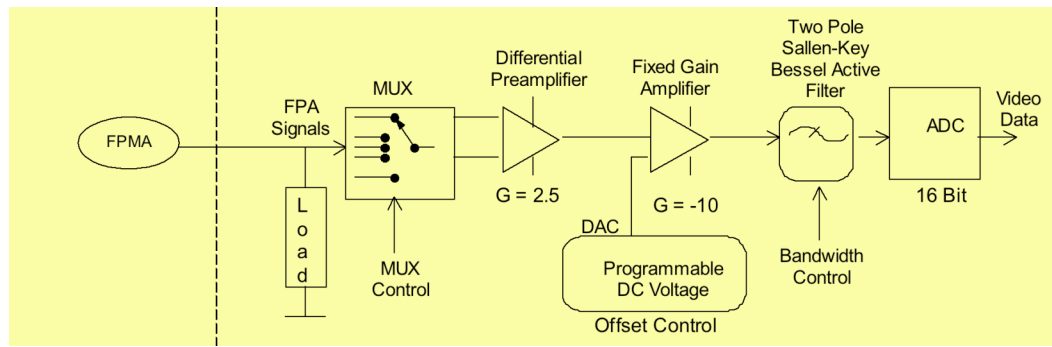
Figure 7.2: Si detector output signal path (one of four parallel paths).

7.1.3.1.2 Array Clock and Operation

There is no provision for accessing only a subset of an array. The Si Analog Signal Processing electronics provide four Si clocks. Observation modes for the IRS require clocking the Si detectors using specific timing patterns as defined below. The Timing Pattern Generator (TPG) is used to clock the Si detectors. Unlike in the MIPS instrument, all the times given here are in standard seconds. In order to discuss the IRS array clocking and data collection it is necessary to define several terms:

Observation: The set of all exposures required to accomplish a single Astronomical Observation Request (AOR).

Sample: A 16-bit integer resulting from a single A/D conversion of a pixel signal.

Frame: A single clocking through a detector array, visiting every pixel. There are four frame types used in observations:

1. ***Sample Frame:*** A frame in which each pixel is sampled (read).
2. ***Reset Frame:*** A frame in which each pixel is reset.
3. ***Boost Frame:*** A frame in which each pixel is reset with a boost voltage applied.

4. ***Spin Frame:*** A frame in which each pixel is just addressed. No read or reset is applied. Spin frames allow continuous clocking patterns without affecting the charge of a pixel.

Frame Count: The number of sample frames.

Two data collection techniques are used in general science data collection: Spinning Double Correlated Sampling (DCS) and Raw or Sample-Up-the-Ramp (SUR) data collection. Science data are collected in SUR mode, while peak-up target acquisition employs DCS.

Spinning Double Correlated Sampling (DCS): Scheme in which the first frame after a reset frame and the frame just before the next reset are sampled, with spins in between.

IRS Raw or Sample-Up-the-Ramp (SUR): Scheme in which sample and spin frames are obtained between a reset frame and the next reset/boost frame; i.e., the ramp is sampled.

In addition, we define different time amounts:

Frame Time: Time required to clock through one Si frame. IRS frame times are in standard seconds (derived from a 40 MHz oscillator on the CE I/O board).

Exposure Time: The time for which the Timing Pattern Generator is running during an IRS exposure ("wall clock time").

Effective Data Collection Event (DCE) Time, or Effective Exposure Time: The time for which the Timing Pattern Generator is actively collecting valid data, or the time between detector resets ("integration time").

DCE Time or Execution Time: The time between and including detector resets (= effective DCE time + reset time).

7.1.3.1.2.1 Double Correlated Sampling

Figure 7.3 shows the parameters for the DCS data collection technique. This is the way in which the peak-up arrays are read for the purposes of IRS Peak-Up. Following an initial series of bias boost and reset frames, each pixel is sampled. The total time required to sample every pixel in the array is the frame time. Several successive spins through the array are then carried out in which each pixel is accessed but not sampled. (Note that the frame time and the spin time are equal). After a specified number of spins, the pixels are again sampled, and the difference between the two samples at each pixel is saved and stored as an image. Each DCS mode sub-exposure will generate a single 128×128×16-bit image.

7.1.3.1.2.2 Sample-Up-the-Ramp Data Collection

Figure 7.4 shows the parameters for the SUR data collection technique. This is

how science exposures (peak-up imaging and spectroscopic exposures) are read. Following an initial series of bias boost and reset frames, each pixel is sampled. The total time required to sample each pixel in the array is the frame time. There may be spin frames in which each pixel is accessed but not sampled. Every time the pixel is sampled, the data from each sample are saved and stored as an image to be sent to the ground for processing. Each IRS SUR mode exposure will generate N, 128×128×16-bit images, where N is the number of commanded sample frames. Note that some combinations of clocking patterns and frame times are inefficient or otherwise undesirable; hence, only a subset of all possible combinations are available to observers (see section 7.2.3).
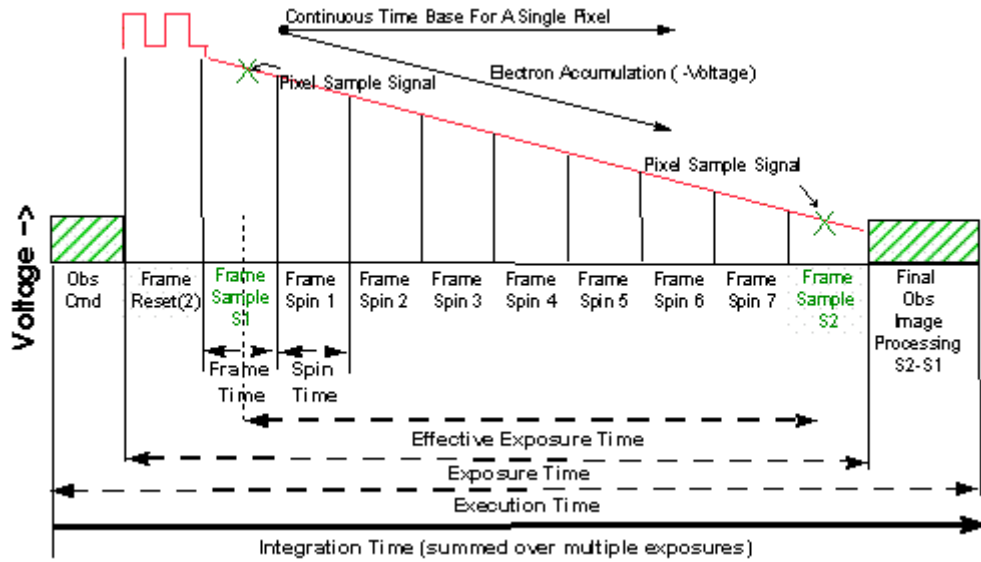
Figure 7.3: Double Correlated Sampling (DCS) data collection technique.
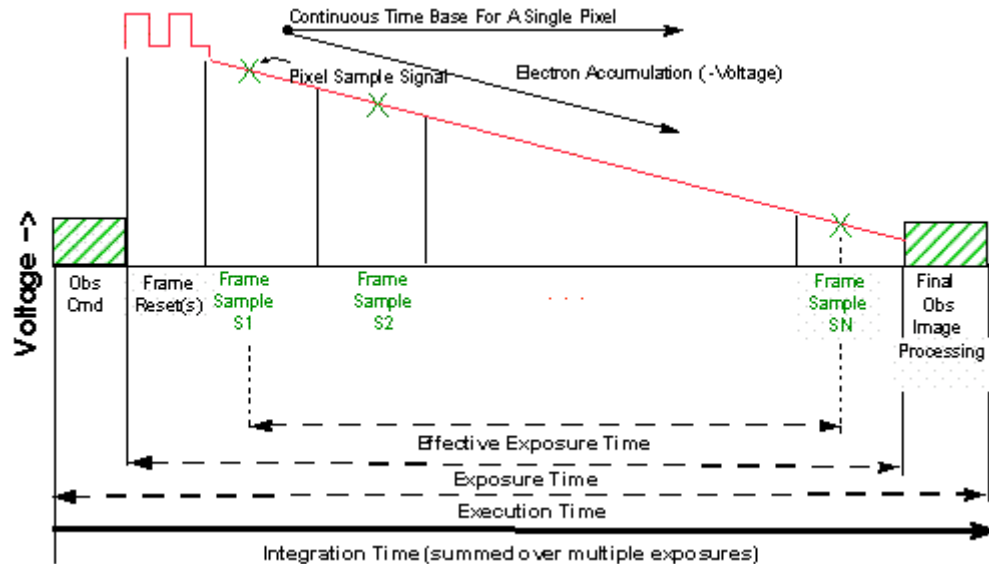
Figure 7.4: "Sample-Up-the-Ramp" data collection technique.

Table 7.3 lists the number of samples ($N_{samples}$) used in SUR data collection mode for each available IRS integration time (t_{int}). The corresponding time interval between samples ($\Delta t_{samples}$) is found by dividing the integration time by the number of samples less one [$t_{int}/(N_{samples}-1)$]. The times given in the table are shorthand for the effective exposure times given in parentheses (see section 7.2.3).

Table 7.3: SUR data collection mode sampling parameters (see text for explanation).

t_{int}		$N_{samples}$	$\Delta t_{samples}$	
6	(6.27)	4	2	(2.095)
14	(14.63)	8	2	(2.095)
30	(31.35)	16	2	(2.095)
60	(60.90)	16	4	(4.065)
120	(121.80)	16	8	(8.125)
240	(241.80)	16	16	(16.125)
480	(481.65)	16	32	(32.115)

7.1.3.1.3   Radiation Effects

The high-energy radiation environment for the IRS consists primarily of (1) cosmic ray protons and heavy ions, and (2) solar protons and heavy ions. These particles range in energy from a few eV to over 1 GeV. The peak fluxes and total fluences of these particles vary with the level of solar activity and are on the order of 2–16 protons cm^{-2} s^{-1}. This environment can cause a variety of effects in the IRS, manifesting as both long-term degradation and abrupt changes in the electronics and optics. The CE is designed to withstand the total radiation dose estimated over a 5-year mission. However, the largest contribution to the accumulated radiation dose damage occurs during unusually large solar flares (see section 7.3.4).

7.1.3.2 Control Section Flight Software

The Control Section Flight Software (CSFS) is the main control and coordination center for the IRS (and MIPS) instruments. Although certain survival commands and engineering telemetry data are passed directly between the instrument and the spacecraft, the major command and telemetry functions are provided by the CSFS. There are no software state or operations mode changes that are selectable by the observer. Upgrades to the software and updates to the IRS software patchable constant table are maintained by the Spitzer Flight Operations and Science Operations teams.

The CSFS provides the following capabilities:
- Receive and process command data.
- Control and coordinate the IRS instrument activities based on defined operational parameters within the accepted commands.
- Configure the IRS detectors and control the collection of science and

calibration image data via the Si array electronics interfaces.
- Perform peak-up target acquisition using the IRS Peak-Up array.
- Gather, format, and output both engineering and science telemetry data.
- Control each of the instrument stimulators.
- Monitor instrument safety parameters and maintain the instrument in a safe configuration.
- Perform memory self-tests.
- Initialize and maintain the instrument's Control Section electronics and software systems.
- Monitor and control each of the instrument array temperatures.
- Control the annealing and radiation mitigation of each of the instrument arrays.

7.1.4 Detectors

Each of the four IRS optical modules has a single Focal Plane Mount Assembly (FPMA) 128×128 BIB (blocked impurity band) focal plane array detector. Two FPMAs contain arsenic doped silicon (Si:As) arrays, and two contain antimony doped silicon (Si:Sb) arrays. The Si:As arrays operate over a 5–26 µm wavelength window, while the Si:Sb arrays operate over a 14–40 µm window. Each array is electrically connected via short cables to the Cold Interface Board (CIB) located within the cold instrument section on the IRS baseplate (see Figure 7.1). As the Combined Electronics is redundant (Side A and Side B), the CIB provides a reliable means of signal interconnection and distribution to each of the four modules.

7.1.4.1 Detective Quantum Efficiency

The Detective Quantum Efficiency (DQE) is a function of wavelength (λ) and bias voltage. In Si:Sb arrays, an increase in bias voltage leads to a dramatic change in DQE at longer wavelengths ($\lambda > 36$ µm); see Figure 7.5. However, it also exacerbates the dark current associated with radiation-damaged pixels. A performance optimization metric, developed from charged particle test data and infrared responsivity vs. wavelength measurements on Si:Sb detectors, is used to calculate the optimum bias voltage for a given cosmic ray dose. In order to reduce the number of damaged pixels ("rogues"), the detector bias for the Long High array was reduced from 2 to 1.6 V. This change happened at the start of IRS Campaign 25 (10 October 2005). The change reduced the detector's quantum efficiency by a wavelength-dependent factor (with an average of ~30%). On the other hand, the change also reduced the number of bad pixels in the array by a factor of two to three, resulting in cleaner spectra.

> **As charged particle damage to the array accumulates over the course of the Spitzer mission, the bias will be reduced to partially offset the associated noise increase.**

7.1.4.2 Array Temperature Control and Annealing

In order to achieve optimum performance, the detector array temperature must be closely controlled. Once an IRS observing cycle has commenced, detector array temperature over the entire observing cycle (~12 hours) is stable to within 0.1 K. Because charged particle radiation damage accumulation on the pixels causes an increase in baseline dark current, periodic annealing is required to reestablish this baseline. Raising the temperature (annealing) of the array to 25 K (±5 K) for a minimum time of 1 minute accomplishes this. Currently, the IRS detectors are annealed as part of the detector check procedure that runs every three to four days of an IRS campaign, including at the beginning after powering up the IRS and at the end just before shutting down. In addition, all IRS arrays are annealed once every 24 hours, prior to taking dark measurements at the start of each Period of Autonomous Operations (PAO).
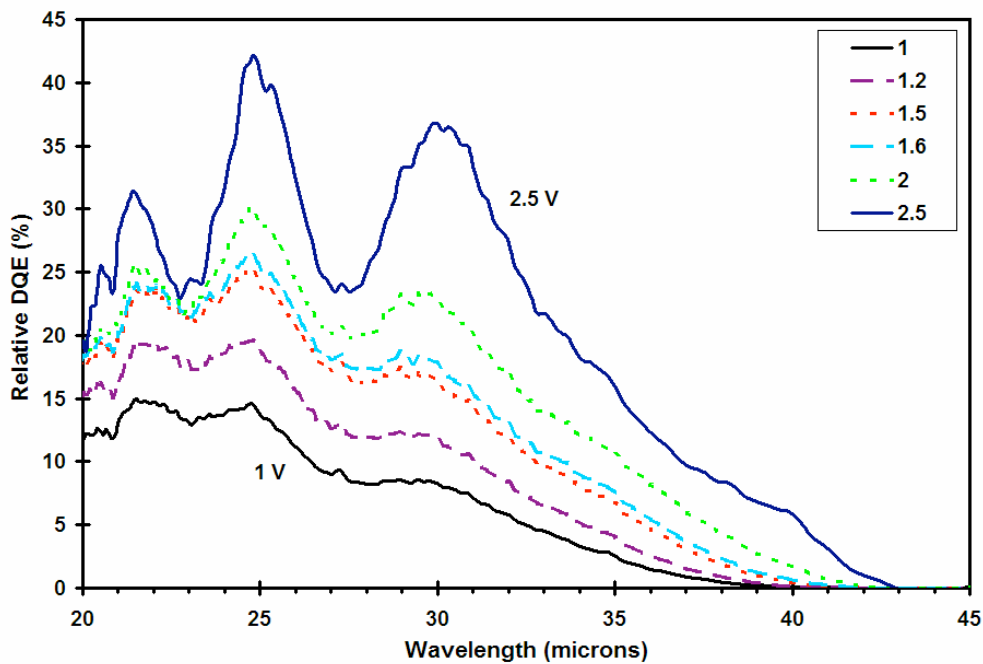
Figure 7.5: DQE for the Si:Sb arrays as a function of wavelength for different bias settings (1–2.5 V). The DQE declines linearly from 30 μm to 40 μm (reaching zero at approximately 40 μm) for the optimal bias voltage (~1.5 V) that provides the best performance over the 14–36 μm interval.

Table 7.4: IRS detector operating temperatures

	Array Temperature Range	Thermal Stability
Si:As	4.7 to 6.2 K	30 mK over a 600 second interval
Si:Sb	3.8 to 5.0 K	20 mK over a 600 second interval

7.1.5 Optics

The IRS optical system covers the wavelength range from 5.2 to 38.0 μm. There is no active focus control for any of the modules. The correct focus of the slit image on the focal plane was achieved by placing an appropriately machined

spacer plate between the Focal Plane Mount Array (FPMA) and the module housing. The spacer plate also corrects for displacements and rotations of the actual detector array relative to its housing. The IRS instrument is capable of self-test for proper functional operation of the focal planes. Through the use of an internal flood illuminating stimulator source, located near the arrays within each of the four Cold Instrument Modules, the detector responsivity can be monitored. All mirrors and gratings are constructed of diamond machined aluminum. The filters are evaporated multi-layer interference filters.

Table 7.5: Sizes of the IRS slits

Module	Wavelength Range (μm)	Pixel Size (arcsec)	Slit width (arcsec)	Slit length (arcsec)
SL 2nd	5.2–8.7	1.8	3.6	57
SL 1st	7.4–14.5	1.8	3.7	57
PU-blue	13.3–18.7	1.8	80	56
PU-red	18.5–26.0	1.8	82	54
LL 2nd	14.0–21.3	5.1	10.5	168
LL 1st	19.5–38.0	5.1	10.7	168
SH	9.9–19.6	2.3	4.7	11.3
LH	18.7–37.2	4.5	11.1	22.3

Since there are no moving parts in the IRS, multiple images/orders are projected onto the detector arrays simultaneously, and telescope motion is used to move a source along the slit. Each aperture in the low-resolution modules is divided into two sub-slits. These sub-slits provide spectroscopy in either the first or second order. Each high-resolution (echelle) module has a single slit. The tolerances of the slit widths are ±5%, with the slit constant in width over its length to ±1%. Table 7.5 summarizes the properties of the IRS slits. Figure 7.6 illustrates the dimensions of the slits and peak-up arrays.

The optical design of each of the IRS spectrograph modules was organized along a set of common design rules. The widths of the entrance slits were set at $f\lambda_{max}$, where λ_{max} is the maximum wavelength for that module and f is the telescope focal ratio ($f = 12$). Internally, each module maps this slit width geometrically across two pixels on its detector array. Slit lengths were set to be as long as possible, limited by the sizes of the detector arrays, the available space in the Cryogenic Telescope Assembly (CTA) focal plane, and the requirement to keep the different orders separated in the cross-dispersed echelle modules. All of the IRS optics are over-sized to minimize internal diffraction effects, and have their optical surfaces coated with gold to enhance their infrared reflectivity. With the exception of the imaging train of the Short-Low module, none of the modules has an internal Lyot stop.
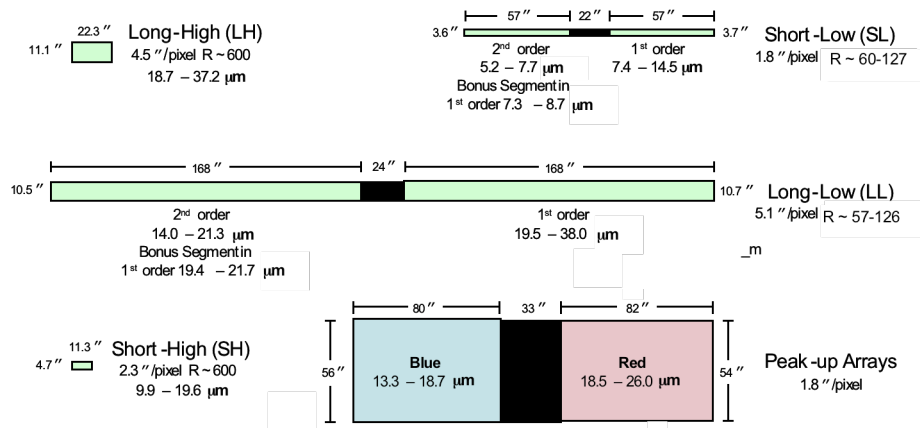
Figure 7.6: Schematic representation of the IRS slits and peak-up apertures. Note that the actual relative orientations and positions in the Spitzer focal plane are shown in Figure 4.5.

> **Note that the IRS slits are *not* parallel in the Spitzer focal plane**; see Figure 2.1 and Figure 4.5 for the relative positions and angles between the slits. The relative angle between the Short-High and Long-High slits is 84.8°, and the angle between the Short-Low and Long-Low slits is 85.9°.

Observers should remember that the interaction of the light beam with the telescope surfaces will result in the appearance of characteristic diffraction patterns at the telescope focal plane (see Figure 4.1). As for other telescopes, these patterns consist primarily of Airy rings (due to the finite telescope aperture) and radial spikes (due to the secondary mirror supports). For bright sources, the power on these features may be significant. For example, experiments with the simulation software STINYTIM (available at http://ssc.spitzer.caltech.edu/archanaly/contributed/browse.html) indicate that, for a source with a brightness of 1 Jy at 8 μm, the extent of the PSF will be detectable at the 1 mJy level up to 10" away from the source. For a source that is 1 Jy at 25 μm, the 1 mJy level is reached at a distance of ~30". Observers that are interested in detecting faint structures next to bright sources should take into account the shape of Spitzer's PSF when designing their AORs. As a rule of thumb, the flux will be smaller than 1/100 of the peak source flux at distances larger than 1.5 to 2 slit widths from the source (for the low-res modules). This rule of thumb is provided as a general guideline only and should not replace a detailed PSF analysis using STINYTIM. See Figure 7.7.
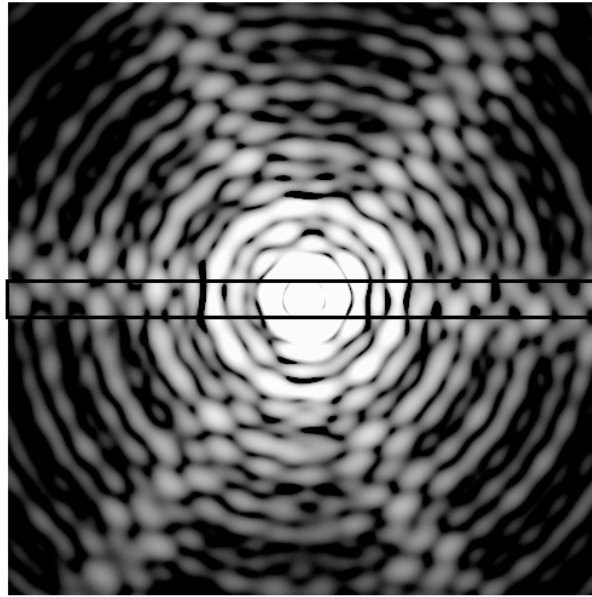
**Figure 7.7: The PSF at 8 μm, perfectly centered in the SL slit (shown as a rectangle 3.6"
wide). Simulated image using STINYTIM.**

7.1.5.1 Short-Low (SL) Module

The Short-Low (SL) module is a grating spectrograph and imager that covers, in
two orders, the nominal spectral range from 5.2 to 14.5 µm at $60 < R < 128$. The
red coverage of the SL module extends to slightly longer wavelength (≈ 15.4 µm)
than indicated here, but SL spectra cannot be calibrated redward of 14.5 µm. The
plate scale for the SL 128×128 Si:As Blocked Impurity Band (BIB) array is
1.8″/pixel. When the source is on one half of the long slit, its light passes through
a 7.5–15 µm bandpass filter and is diffracted by the grating in first order (SL1).
When the source is moved to the other half of the same slit, its light passes
through a 5–7.5 µm bandpass filter and is diffracted by the grating in second order
(SL2). The module also produces a second, short spectral segment that covers the
7.3–8.7 µm spectral region in first order (R=60-72), when the source is falling on
the 2nd order (SL2) sub-slit. This overlap region, known as the "bonus segment,"
is intended to aid the observer in normalizing the first- and second-order segments
of the spectrum. Figure 7.8 shows the optical layout of the SL module.

In addition to obtaining spectra, the SL module provides the IRS peak-up
function. **The purpose of peak-up is to fine tune the telescope's pointing. In
addition, if the science target is used for peak-up acquisition, this will help to
accurately position it onto any of the IRS slits.** Light passes through one of two
bandpass filters (see Table 7.5) and is imaged onto the SL detector array focal
plane (in a different detector region from that used for the spectrum). A
description of the peak-up procedure is given in section 7.2.4. The peak-up array
can also be used to produce small-field mid-IR images; this mode is called Peak-
Up Imaging.
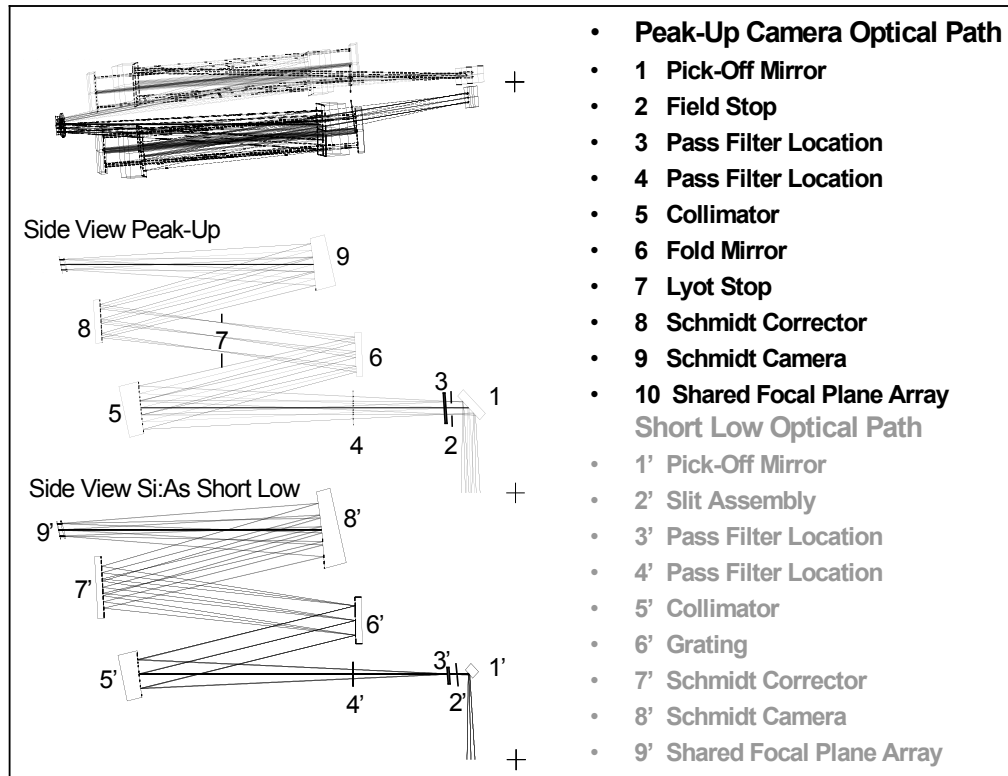
Figure 7.8: Short-Low spectral (grey) & Peak-Up (black) optical components and light paths.

7.1.5.2 Long-Low (LL) Module

The Long-Low (LL) module is a grating spectrograph that covers, in two orders, the nominal spectral range from 14.0–38.0 µm at $57 < R < 126$. The red coverage of the LL module extends to slightly longer wavelength (≈ 42 µm) than indicated here, but LL spectra cannot be calibrated redward of 38.0 µm because of light leakage at the long wavelength end caused by delamination of the order-sorting filter (see below). The detector is a 128×128 Si:Sb BIB array with a plate scale of 5.1″/pixel. As with the SL module, the LL aperture consists of two, in-line sub-slits with a small gap between them. When the source is in one sub-slit, light passes through a 14–21 µm bandpass filter and is diffracted by the grating in second order (LL2). When the source is moved to the other sub-slit, light passes through a 21–40 µm bandpass filter and is diffracted by the grating in first order (LL1). Long-Low also has a "bonus segment," similar to that in SL, which covers 19.4–21.7 µm (R=57–63). Figure 7.9 depicts the LL module components and optical paths.
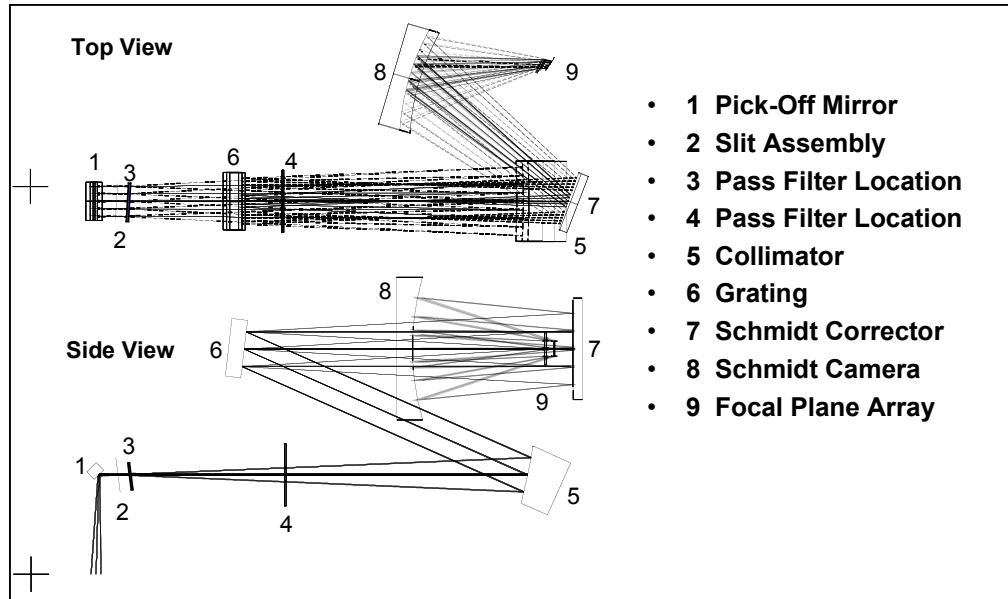
Figure 7.9: Long-Low module optical components and paths.

7.1.5.2.1   *Light Leakage in the Long-Low Module*

The IRS LL module displayed degraded sensitivity in 1st order during CTA ground testing prior to launch due to delamination of the order-sorting filter. The LL1 filter survived launch and is now returning useable science data. Data obtained during the In-Orbit-Checkout (IOC) and Science Verification (SV) indicate that contamination of the LL1 spectrum from filter leakage is no more than 5% (shortward of about 38 μm), based upon observations of hot stars; that is, the flux of the obtained spectrum deviates by less than 5% from the expected spectrum. This leakage is most pronounced beyond about 38 μm; the spectrum can be severely contaminated by blue light at wavelengths longer than about 38 μm.

7.1.5.3 Short-High (SH) Module

The Short-High (SH) module is a cross-dispersed echelle spectrograph that covers, with a single pointing, the nominal spectral range from 9.9–19.6 μm at R ≈ 600. It has a small wavelength overlap (18.7–19.6 μm) with the Long-High module. The detector is a 128×128 pixel Si:As BIB array with a plate scale of 2.3″/pixel. The SH aperture is a single slit of size 2×5 pixels. Its length is limited so that nearly an octave of a cross-dispersed echelle spectrum can be fit onto the detector array without spatial overlap of orders. Figure 7.10 depicts the module components and optical paths.

7.1.5.4 Long-High (LH) Module

The Long-High module is a cross-dispersed echelle spectrograph which covers, with a single pointing, the nominal spectral range from 18.7–37.2 μm at R ≈ 600. It has a small wavelength overlap (18.7–19.6 μm) with the Short-High module.

The detector is a 128×128 pixel Si:As BIB array with a plate scale of 4.5″/pixel. As with the SH module, the LH aperture is a single slit. Its length is limited so that nearly an octave of a cross-dispersed echelle spectrum can be fit onto the detector array without spatial overlap of orders. Figure 7.11 depicts the LH components and optical paths.
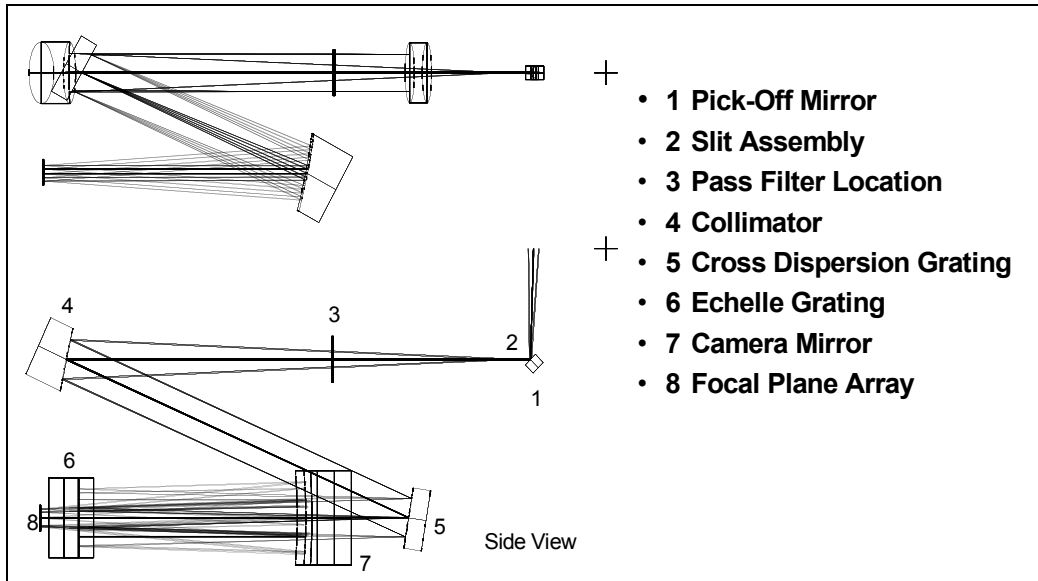
Figure 7.10: Short-High module optical components and paths.
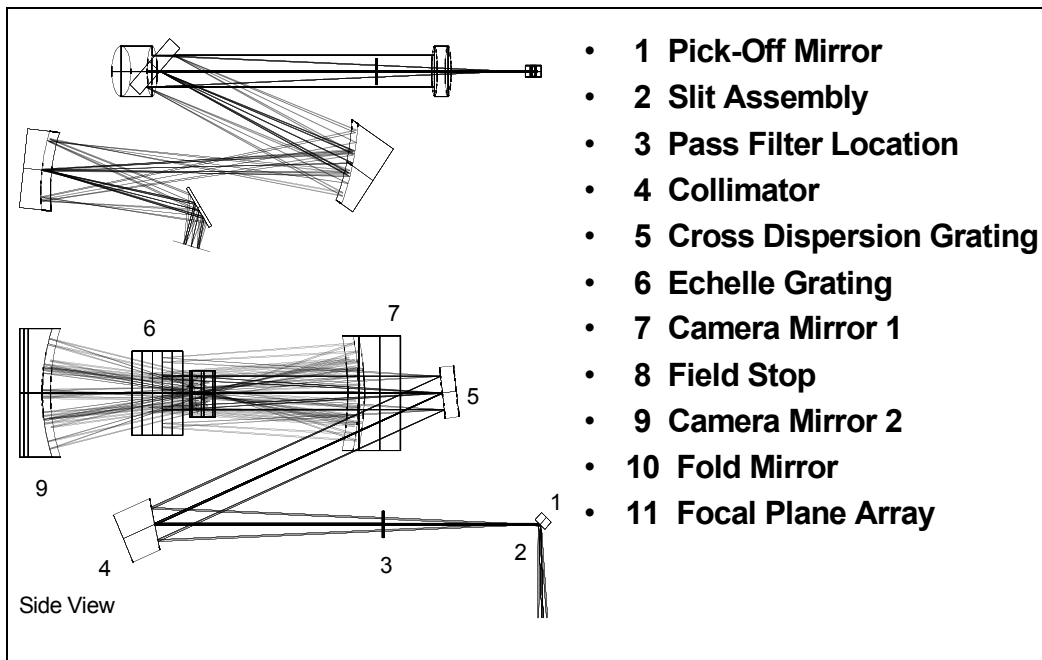
Figure 7.11: Long-High module optical components and paths.

7.1.6 Spectrum and Image Quality

This section shows "typical" IRS data obtained during in-orbit testing. Figure 7.12 shows flat field images for SL (left) and LL (right). These images illustrate the full areas of the detector arrays dedicated to each spectral order. For both modules, the 1st order (longer wavelength) spectrum is on the left and the 2nd order (shorter wavelength) spectrum is on the right; wavelength increases from top to bottom within each order. The short bonus segment is visible above the 2nd order spectrum for both modules. (See Table 7.5 for the wavelength ranges covered by each order.) The SL image also shows the two peak-up arrays. Spectral fringing (see section 7.1.5.2.1) is visible in the LL1 spectrum.
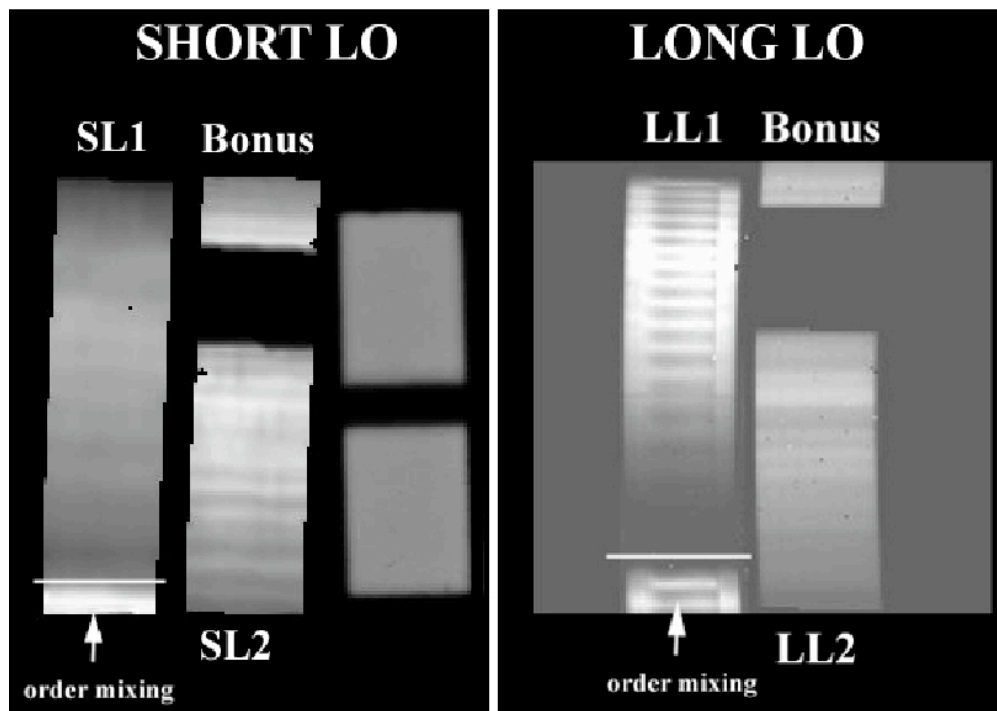
Figure 7.12: Flat field images for the Short-Low (left) and Long-Low (right) modules.
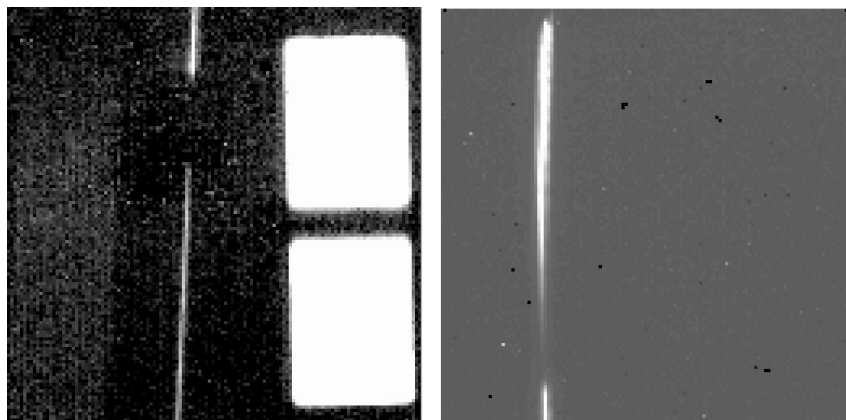
Figure 7.13: Unprocessed data from the SL (left) and LL (right) modules obtained in orbit for a point source.

Figure 7.13 shows an unprocessed 2-D image from each of the SL (left) and LL (right) modules. These images are the result of observations of a point source as part of IRS tests conducted during IOC and SV. The SL image shows a spectrum obtained in 2nd order (along with the 1st order bonus segment at the top and the saturated IRS peak-up arrays on the right), while the LL image shows a spectrum obtained in 1st order.

The low resolution spectrographs have the following resolving power (λ in μm, 1σ ~5 in R):

- SL1, SL bonus: R=8.2667×λ
- SL2: R=16.5333×λ
- LL1, LL bonus: R=2.9524×λ
- LL2: R=5.9048×λ .

Figure 7.14 shows an unprocessed 2-D image from each of the SH (left) and LH (right) modules. As with the low-resolution data shown above, these images were obtained during IOC testing and are spectra of a point source. Both images show the typical layout of a cross-dispersed echellogram. For SH, wavelength decreases from left to right (the blue orders are closer together than the red) and within each order, wavelength decreases from bottom to top. For LH, wavelength increases from left to right with diffraction order and increases from bottom to top within each diffraction order. **All of the orders and slits in a given IRS module are read together in one complete read cycle (this is also true of the low-resolution modules). Only one module (array) is operational at any given time.** The high resolution spectrographs have R=600. This has been confirmed in-orbit within 1σ =60. The large uncertainty reflects errors induced by line undersampling for in-orbit measurements.

Figure 7.15 shows a sample in-orbit image from the SL array with a point source (i.e., a star) centered in the blue peak-up sub-array. Note that the first diffraction ring of the point source is clearly visible. Bright pixels in the image are cosmic ray hits. The illuminated red peak-up sub-array is visible in the figure above the blue sub-array, but does not contain any stars. For a description of the IRS peak-up algorithm, see section 7.2.4.

Figure 7.16 and Figure 7.17 show examples of fully extracted, processed, and calibrated IRS 1D spectra of astronomical objects. Figure 7.16 shows a low-resolution spectrum of the ultraluminous infrared galaxy UGC 5101 that utilizes both SL and LL data, with exposure times of 12 seconds and 28 seconds in each module subslit, respectively. Figure 7.17 shows a high-resolution spectrum of the starburst galaxy NGC 7714 that utilizes both SH and LH data, with exposure times of 240 seconds in each module. These spectra are presented in Houck et al. (2004, ApJS, 154, 18), and each is discussed in detail in Armus et al. (2004, ApJS, 154, 178) and Brandl et al. (2004, ApJS, 154, 188), respectively.
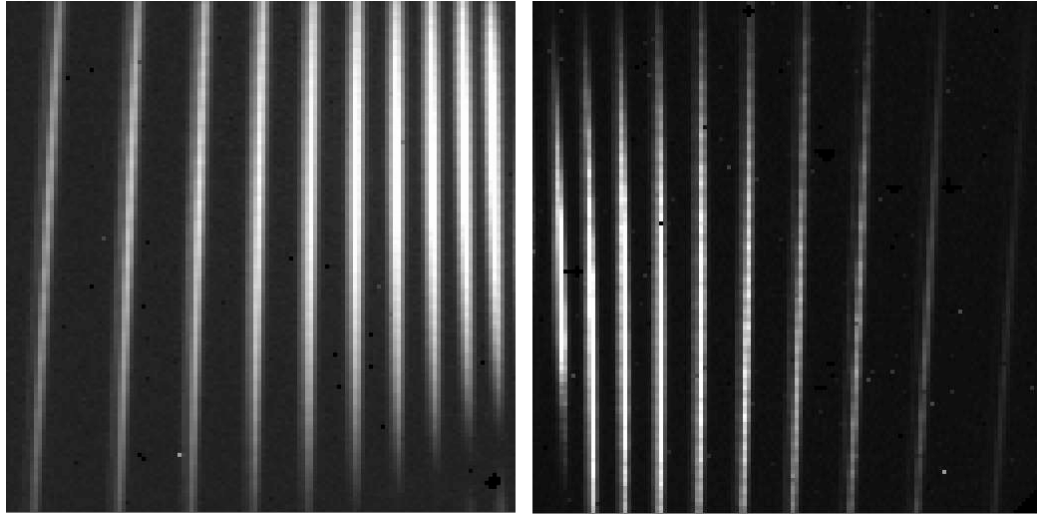
Figure 7.14: Unprocessed data from the SH (left) and LH (right) modules obtained in orbit for a point source.
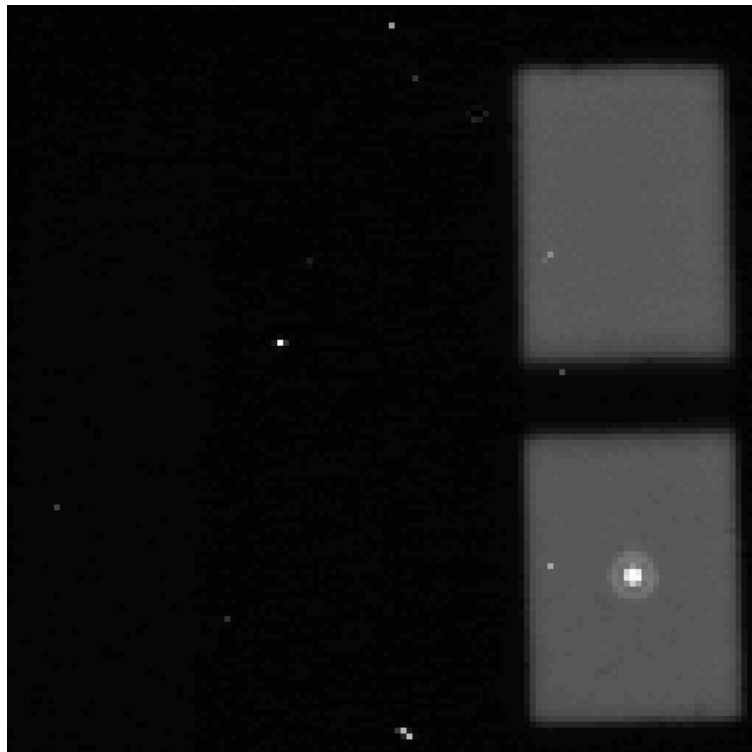
Figure 7.15: In-orbit image of a point source on the blue IRS peak-up sub-array.
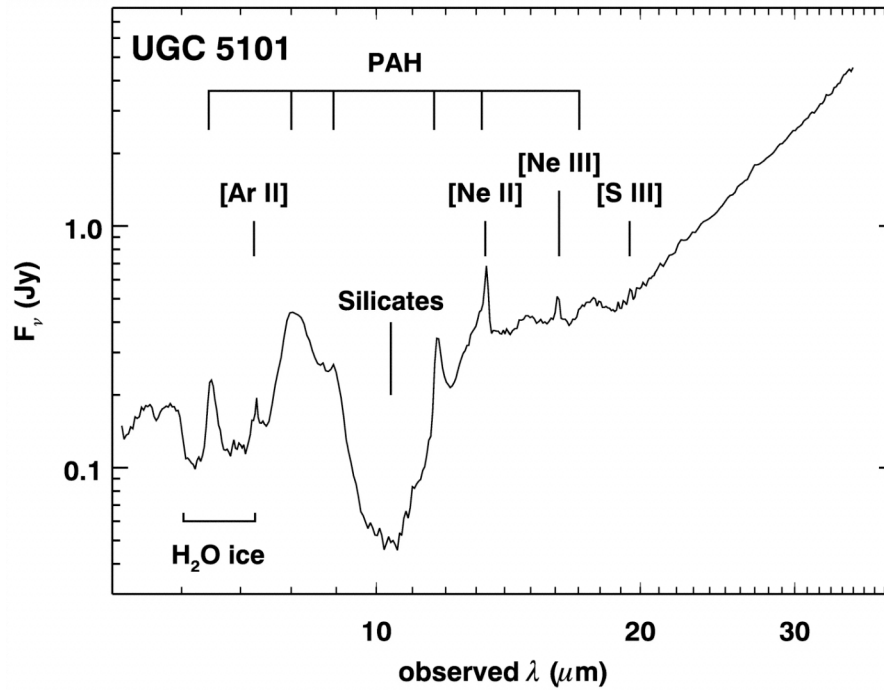
Figure 7.16: IRS low-resolution (SL + LL) spectrum of UGC 5101 (from Houck et al. 2004, ApJS, 154, 18 and Armus et al. 2004, ApJS, 154, 178). Several strong absorption bands from PAHs, dust, and ice, as well as atomic emission lines, are visible (redshifted for z = 0.039).
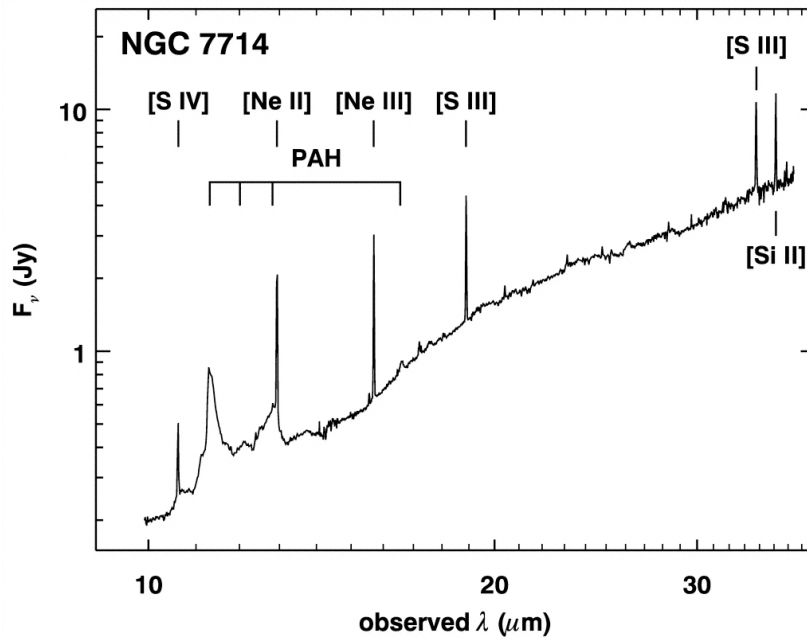
Figure 7.17: IRS high-resolution (SH + LH) spectrum of NGC 7714 (from Houck et al. 2004, ApJS, 154, 18 and Brandl et al. 2004, ApJS, 154, 188). Prominent spectral features are visible as in Figure 7.16 (redshifted for z = 0.0093).

7.1.7 IRS Performance Characteristics

The following reflects the current knowledge of the in-orbit performance of the IRS. Updates to this information will be posted on the SSC website.

1. ***Photometric response*** – Tests have shown that the detector response stability is better than 1% over multiple cool-downs separated by months, and that the responsivity of the detectors are unaffected by radiation. Overall repeatability depends upon the selected pointing (peak-up) accuracy but can be as good as 2%.
2. ***Effect of overexposures*** – Collection of over 300,000 electrons per pixel before a reset can cause a 1–2% reduction in responsivity in the affected pixels for a period of approximately one hour.
3. ***Read noise*** – The read noise is 30 electrons per pixel in SUR ("sample up the ramp") mode with 16 samples per ramp (see Figure 7.4).
4. ***Dark current*** – At a bias of 1.5 volts, the dark current is less than 10 electrons per second per pixel for the Si:As arrays and less than 40 electrons per second per pixel for the Si:Sb arrays.
5. ***Gain*** – The gain is 4.6 electrons per DN (data number).
6. ***Pixel masks*** – The pixel badness criteria are (1) having a dark current greater than 40 electrons per second in Si:As or greater than 160 electrons per second in Si:Sb, or (2) a responsivity less than 50% (or greater than 150%) of the median responsivity of the array.
7. ***Dark frames*** – Since the IRS has no internal shutter, reference frames are not truly dark. Reference spectra are collected at the ecliptic poles, or near the target if the local background is significantly different from the high latitude zodiacal sky. However, each array has areas that are not directly illuminated by the optics. These areas are also used to monitor dark current performance.
8. ***Flat fields*** – For the low resolution modules, cross-dispersed response calibrations are obtained from measurements of the zodiacal emission at low ecliptic latitudes, and spectral corrections are obtained by scanning standard stars along the slits. For the high resolution modules, both dimensions of the flat field frames are obtained from the stellar sources.
9. ***Spectral resolution and wavelength calibration*** – Spectral resolution of the low resolution modules is $R=8.2667\lambda$ for SL1 and SL bonus, $R=16.5333\lambda$ for SL2, $R=2.9524\lambda$ for LL1 and LL bonus, and $R=5.9048\lambda$ for LL2 (for λ in μm, 1σ R~5). Both of the high resolution modules exhibit fixed resolutions, measured to be 600±60, with the uncertainty reflecting line undersampling. The wavelength calibration is based entirely on flight data, and has an accuracy of 1/5 resolution element, limited by the undersampling.
10. ***Cosmic ray transient effects*** – Proton irradiation shows no cosmic ray transient effects after the usual boost/reset sequence (see section 7.1.3.1.2). However, there may be transient effects within a given ramp for a small number of reads before the next boost/reset — this is corrected in the science data processing pipeline.
11. ***Background*** – Typical ecliptic pole photocurrents for the SL, SH, LL, and LH modules are 14, 3, 70, and 8 electrons per second per pixel, respectively, at

nominal operating biases.

12. **Bias** – Nominal operating bias will be 2 volts for SL and SH, 1.8 volts for LL, and 1.6 volts for LH. Before IRS Campaign 25 (10 October 2005), bias voltage was 2 volts for LH. As part of the normal maintenance plan for the IRS, the bias level will occasionally be adjusted.

13. **Latent images** – Latent images are suppressed to less than the 1% level between science ramps by the detector clocking patterns, and by keeping the detector as warm as possible without increasing the dark current. Post-launch measurements in the SL module indicate that latent images amount to between 0.04 and 0.1% of the flux in the previous image and decay exponentially, with an e-folding time of ~1100 sec. Both the magnitude and the decay time depend on the number of collected photons per pixel per ramp, and these values apply strictly only to the range $10^4 - 5 \times 10^4$ e- pix^{-1}.

14. **Droop** – Droop is a slope proportional to the photocurrent summed over all pixels, which is added to every pixel of the array when it is non-destructively read in SUR mode. Noise measurements show that this slope is a multiplexer (MUX) artifact, not a true current.   The droop coupling constant (droop/[total array photocurrent]) is typically 3×10^{-5} for all arrays, which is equivalent to a ratio of droop to average array current of 0.48. Droop is present in all IRS arrays.

15. **Well depth** – Well depth, defined as the point at which the differential responsivity has fallen to 90% of its small-signal value, exceeds 2×10^5 electrons per pixel for Si:As and exceeds 10^5 electrons per pixel for Si:Sb. The full well depth is a few times these values and irrelevant for the user, because digital saturation (301,466 electrons) is reached before electronic saturation.

7.2 How to Use the IRS

7.2.1 Performance of the Instrument

A number of factors affect the expected performance of the IRS and must be considered when planning observations. Notable among these are sensitivity and signal-to-noise (S/N) estimates, target parameters (e.g., point vs. extended sources, brightness), background brightness, and background subtraction planning (which is important when the science target is in close proximity to a bright mid-IR source). These topics are discussed in the following sections.

7.2.1.1 Sensitivity and Signal-To-Noise Calculations

This section is intended to illustrate the procedure by which the IRS Instrument Support Team calculated the expected instrumental sensitivity and signal-to-noise values. Observers are urged to use the sensitivity plots shown in section 7.2.1.2 or the Spectroscopy Performance Estimation Tool (SPEC-PET) available on the SSC website for observation planning.

We calculate the noise-equivalent flux density (colloquially referred to as "sensitivity") for the IRS via the following steps:

1. The background photocurrent from sky and telescope is given by

$$i_{sky} = (\lambda B_\lambda / hc) A\Omega\Delta\lambda\tau\eta G \tag{7.1}$$

where h is the Planck constant, and

$$B_\lambda = f_{stray} 2hc^2 (1/\lambda)^5 10^8 \Sigma\varepsilon_i / e^{14387/\lambda T_i} - 1) \; [\text{Jy sr}^{-1}] \tag{7.2}$$

where $2hc^2 = 1.19 \times 10^{-4}$ W μm^2, all lengths are in μm, f_{stray} = the ratio of total sky throughput to that of an ideal f/12 telescope, and ε_i are sky emissivities at temperatures T_i. We also define the following:

$A\Omega$ = the telescope throughput in cm² sr.
$\Delta\lambda$ = the spectral bandpass falling on one pixel ($=\lambda/R$) in cm. Note: while there are 2 pixels per resolution element, the slit width is also two pixels. Hence, $\Delta\lambda$ is proportional to the slit width if the spectral plate scale is unchanged.
T_i = effective blackbody temperature of sky + telescope.
τ = the cold optical throughput for extended sources (i.e., omitting slit losses).
ηG = the detector responsivity (electrons per photon) $[=(\eta/\beta)\beta G]$.
η = quantum efficiency of the detector at the desired wavelength.
βG = the dispersion-gain product.

2. In the faint source limit (see sections 7.2.1.2 and 7.2.1.2.2), the total noise in electrons per pixel is the sum of the photocurrent shot noise (i_{sky}), read noise, and dark current shot noise; it is given by:

$$N = \sqrt{i_{sky}\beta G t_{int} + (RN)^2 + i_d \beta_d G_d t_{int}} \; [\text{e- px}^{-1}] \tag{7.3}$$

where i_d is the dark current contribution, RN is the read noise in electrons, t_{int} is the integration time in seconds, and $\beta_d G_d$ is the dispersion-gain product for the dark observations.

3. The relationship between photocurrent per pixel and the incident flux density of a source in mJy is given by:

$$S_{mJy} = 1.511\pi(D^2/4)\tau_{tel}\tau_{align}\tau\tau_{slit}\eta G t_{int} / 4R \; [\text{e- px}^{-1} \text{ mJy}^{-1}] \tag{7.4}$$

where 1 mJy = 10^{-26} erg s^{-1} cm^{-2} Hz^{-1}, D = the telescope aperture (cm²), τ = the

slit throughput of a point source with angular diameter λ_{max} / D (without considering losses) and $\tau_{tel}, \tau_{align}, \tau_{slit}$ are throughputs accounting for losses in the telescope, alignment effects, and slit, respectively. The second factor of 4 in Equation 7.4 comes from dividing the point source flux over 2 pixels per spectral resolution element and 2 pixels spatial extent.

4. Again assuming the faint source limit (see sections 7.2.1.2 and 7.2.1.2.2), the 1σ staring point source continuum sensitivity (PSSC) in mJy at full resolution of the spectrograph is:

$$PSSC = \alpha N / S_{mJy} \quad [\text{mJy}] \tag{7.5}$$

where the factor α accounts for the variation in PSF with wavelength and details of the point source extraction; its value lies between 0.8 and 1.2. The 1σ staring point source continuum sensitivity (in mJy) smoothed at the required resolution of the spectrograph ($R_{req} = 50$) is referred to as the *PSSCS*, and is given by

$$PSSCS = PSSC\sqrt{R_{req} / R} . \tag{7.6}$$

The spectra are smoothed from R to R_{req} and the S/N is assumed to increase as the square root of the number of pixels averaged; suitable sub-pixel smoothing is assumed. This is the quantity given in the sensitivity curves shown in section 7.2.1.2.

5. The 1σ staring point source sensitivity, for unresolved lines (in W m^{-2}) is then

$$PSSL = 3\times10^{-15} \, PSSC/(R\lambda) \tag{7.7}$$

where the factor $3\times10^{-15} = c$ (μm/s) (10^{-29} W m^{-2} Hz^{-1}/mJy), *PSSC* is in units of mJy, and λ is in μm.

7.2.1.2 Sensitivities of the IRS Modules

This section provides sensitivity curves for all four IRS modules. These sensitivities are based on pre-launch models that have been refined using in-orbit measurements of point source flux calibrators, and take into account the degradation of the LL1 module that was discovered during CTA ground-testing prior to launch (see section 7.1.5.2.1). Figure 7.18 shows a comparison between the four modules. Plots of the point-source single pointing continuum sensitivity of the SL and LL modules are shown in Figure 7.19 and Figure 7.20, respectively. Plots of the point-source single pointing (unresolved) line sensitivity of the SH and LH modules are shown in Figure 7.21 and Figure 7.22, respectively. These plots also show the sensitivities for the four observer-selectable integration times. See section 7.2.3.4.4 for a discussion of sensitivities in the Peak-Up Imaging mode. Finally, Figure 7.23 shows the IRS digital saturation limits for bright point

sources based on in-orbit measurements with Spitzer for the shortest ramp time (6 sec). Saturation is not necessarily fatal for an observation as the data processing pipeline (section 7.3) will be able to reconstruct the ramp if at least two samples are available (albeit with a larger error than usual). The SSC recommends that the observer designs his/her observations using multiple cycles of short unsaturated ramps.

However, very bright sources may leave latent imprints in the detectors that will compromise further observations until an anneal is performed (section 7.1.4.2). *The SSC reserves the right not to observe some targets as a result of impacts these bright objects might have on subsequent observations. Observers planning observations of bright objects should discuss the impact of such observations in the proposal.*

Experience with IRS data has shown that residual spectral fringes may remain even after flat-fielding during the standard pipeline processing. While this is most likely to occur in SH and LH data, fringes can also be present in LL data.

Signal-to-noise estimates based on the sensitivity curves in the SOM (which are also available on the SSC website and are used in the online SPEC-PET tool) assume an "ideal" spectrum that is not affected by signal degradation caused by extraction, cosmic ray damage to the detectors, and other factors, such as errors in the flatfield or fringes in the spectrum. The maximum achievable signal-to-noise values are currently ~100 for the low resolution modules and ~50 for the high resolution ones.

In all of the figures shown in this section, it is assumed that the source is in the faint source limit (i.e., for a selected exposure time, the target flux is BELOW the relevant line shown in Figure 7.24). If the source is bright and lies above the line, then the curves shown in the following figures will significantly **OVERESTIMATE** the expected signal-to-noise ratio. A simple scaling rule is given in section 7.2.1.2.2 that can be used to convert the expected S/N ratio based on the (faint source) sensitivities to a more realistic value for bright sources.

In order to reduce the number of unreliable pixels (rogues) in the LH detector, the bias voltage was changed from 2 to 1.6 volts in 10 October 2005. This reduced the sensitivity of the detector by ~30%. On the other hand, the number of "rogue" pixels was reduced by a factor of 2 to 3, resulting in cleaner spectra.

As part of the normal maintenance plan for the IRS, the detector bias levels will occasionally be adjusted. Each change in the bias level will change the module sensitivities, and will affect the long wavelength modules more than the short wavelength modules. Refer to the SSC website for the most up-to-date IRS sensitivity information.
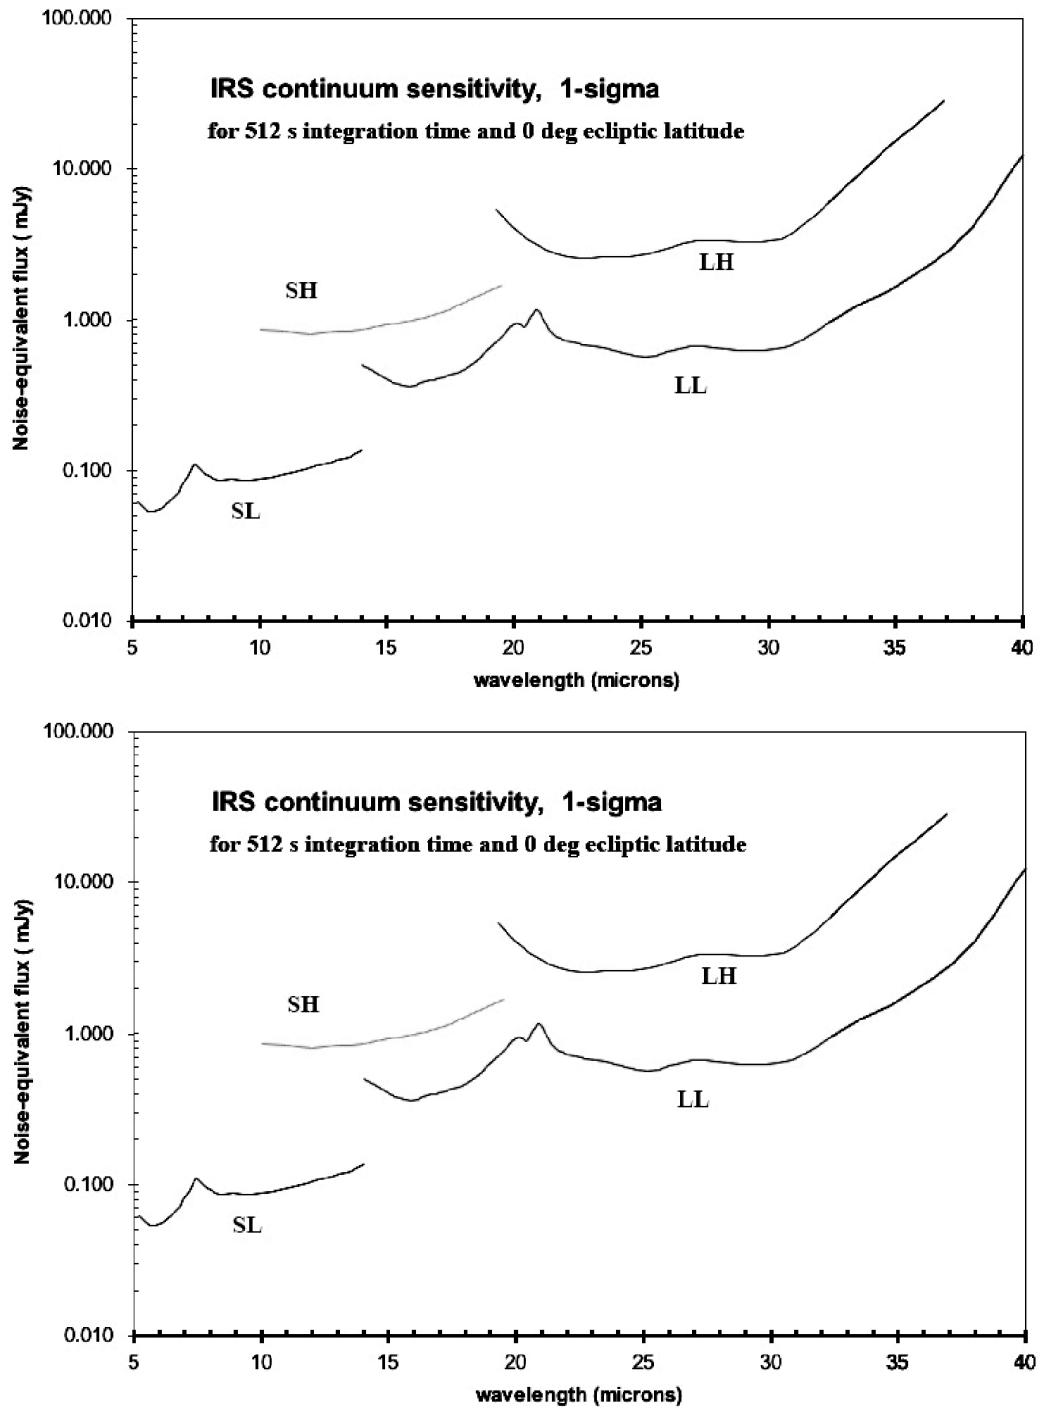
Figure 7.18: 1σ **continuum sensitivity (in the faint source limit) of all four IRS modules in a 512 second integration with HIGH (top) and MEDIUM (bottom) backgrounds (the LOW background case, corresponding to an ecliptic latitude of 90°, is indistinguishable from the MEDIUM background case at the scale of these plots).**
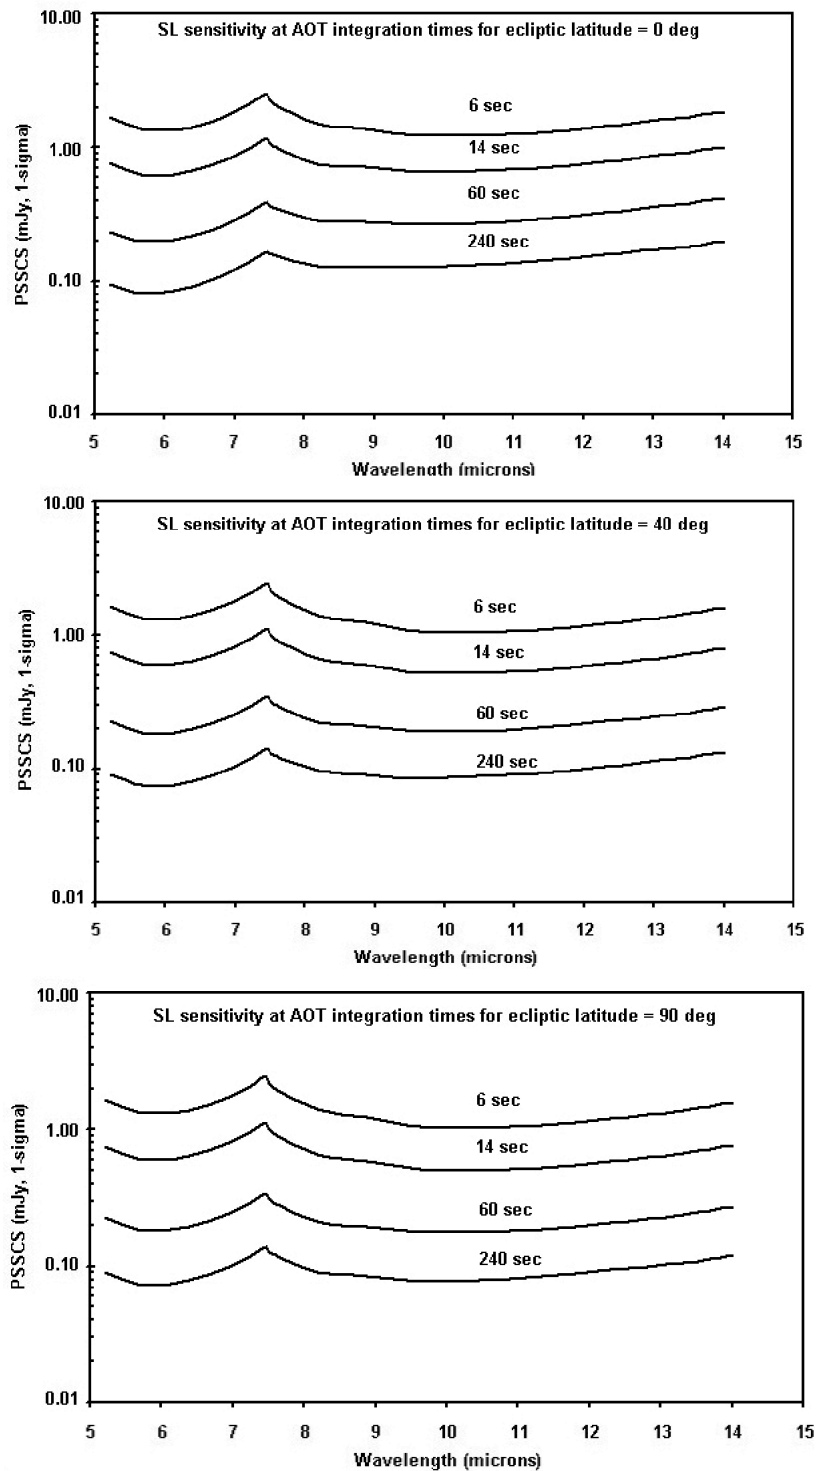
Figure 7.19: Short-Low 1σ point source staring continuum sensitivity (PSSCS) in the faint source limit after smoothing to a resolution of R=50, with HIGH (top), MEDIUM (middle), and LOW (bottom) backgrounds.
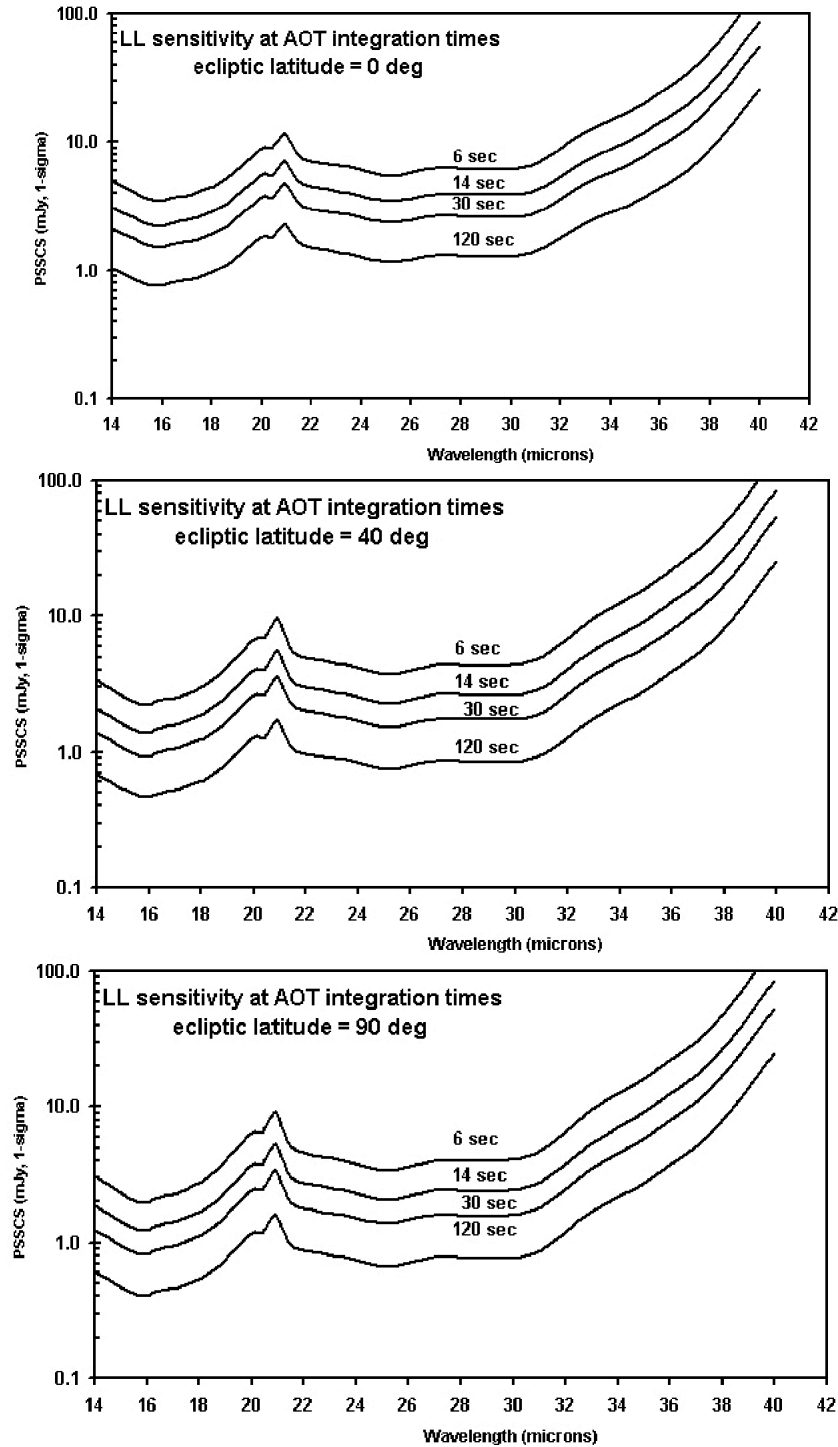
Figure 7.20: Long-Low 1σ point source staring continuum sensitivity (PSSCS) in the faint source limit after smoothing to a resolution of R=50, with HIGH (top), MEDIUM (middle), and LOW (bottom) backgrounds. This figure includes the effects of the LL1 filter delamination discovered before launch.
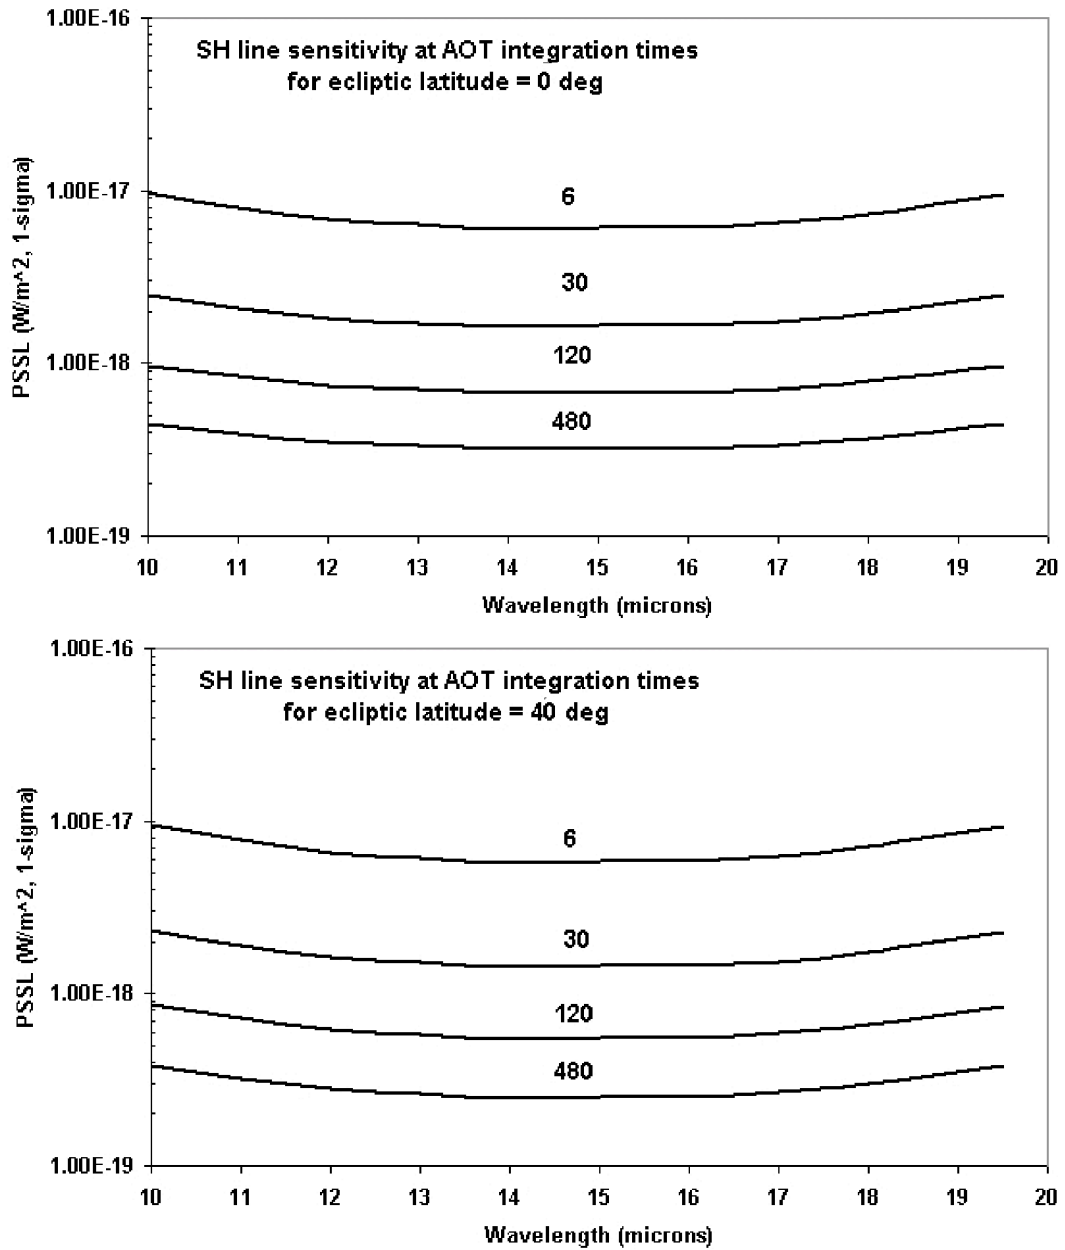
Figure 7.21: Short-High 1σ point source sensitivity (PSSL) to an unresolved emission line in the faint source limit, with HIGH (top) and MEDIUM (bottom) backgrounds (the LOW background case, corresponding to an ecliptic latitude of 90°, is indistinguishable from the MEDIUM background case at the scale of these plots).
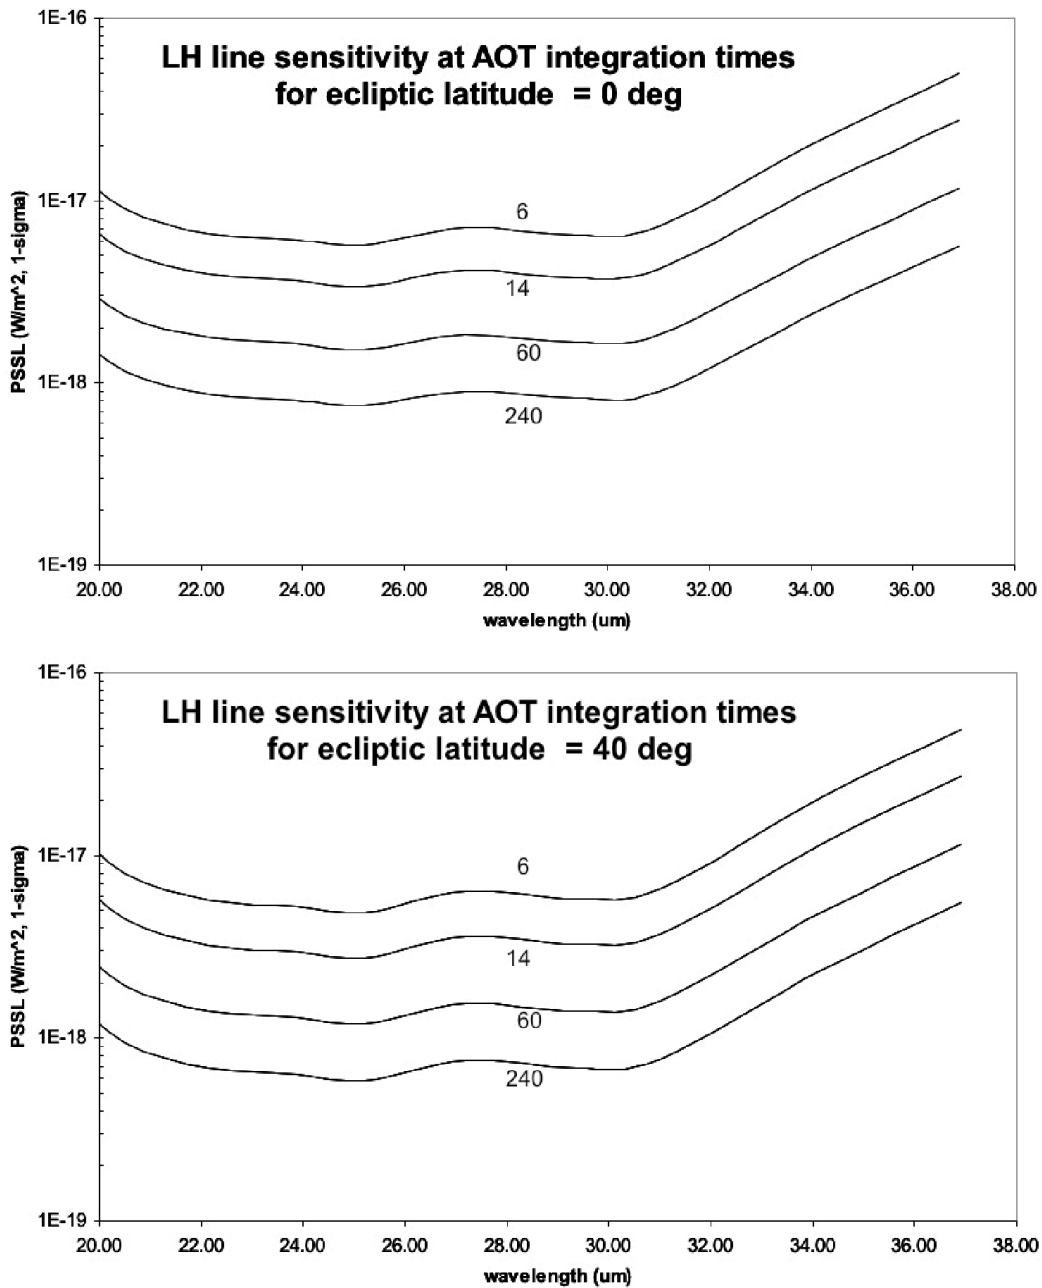
Figure 7.22: Long-High 1σ point source sensitivity (PSSL) to an unresolved emission line in the faint source limit with HIGH (top) and MEDIUM (bottom) backgrounds (the LOW background case, corresponding to an ecliptic latitude of 90 degrees, is indistinguishable from the MEDIUM background case at the scale of these plots).
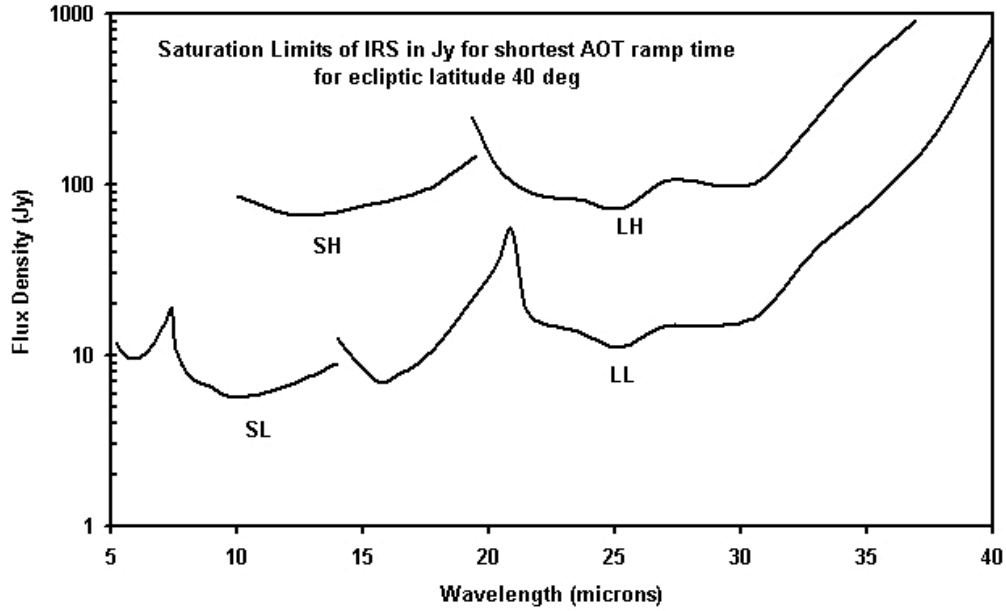
Figure 7.23: Saturation levels for the shortest ramp duration available in the IRS AOTs, with MEDIUM background (the HIGH and LOW background cases, corresponding to ecliptic latitudes of 0 and 90°, respectively, are indistinguishable from the MEDIUM background case at the scale of this plot). For flux densities at a given wavelength that exceed these limits, at least one of the samples will be saturated. Flux densities up to 3 times greater will have at least two unsaturated samples. The observer must decide whether the science goals can be accomplished if the spectrum is saturated in some spectral or spatial regions, and not in others.

7.2.1.2.1 *Sensitivity to Extended Sources*

The calculations shown in section 7.2.1.1 (as well as the sensitivity curves shown in section 7.2.1.2) apply only to observations of point sources. The sensitivity of the IRS to extended sources is currently being investigated as data are accumulated in orbit. For the moment, we provide a rough guide to the conversion from point- to extended-source sensitivities as follows. Note that the approximations used here are only good to about 20% because of wavelength-dependence in the conversion factors that varies from module to module. The ESSC and ESSL are the 1σ staring extended source sensitivities per pixel for continuum and line, respectively. They are:

$$ESSC \sim 8 \times PSSC / \phi^2 \quad [\text{MJy sr}^{-1}] \qquad (7.8)$$

and

$$ESSL \sim 8 \times 10^9 \ (PSSL) / \phi^2 \quad [\text{W m}^{-2} \text{ sr}^{-1}] \qquad (7.9)$$

where ϕ is the linear edge size of a pixel projected onto the sky in arcseconds (determined from column 3 of Table 7.5, and different for each IRS module). The values of *PSSC* (or its smoothed equivalent for a variety of AOT integration

times) and *PSSL* can be read from the sensitivity curves in section 7.2.1.2.

For example, the 1σ point source sensitivity (PSSC) of the SL module at 8 μm is 1.05 mJy in 6 seconds (from the sensitivity curve in Figure 7.19). The pixel size for the SL module is 1.8 arcseconds (from Table 7.5). So, the corresponding extended source sensitivity for SL is given by Equation 7.8 as follows:

$$ESSC(SL) \sim 8 \times 1.05 / 1.8^2 = 2.6 \text{ MJy sr}^{-1} \text{ at 8 μm in 6 sec.} \qquad (7.10)$$

Similar calculations can be carried out for the other IRS modules and exposure times.
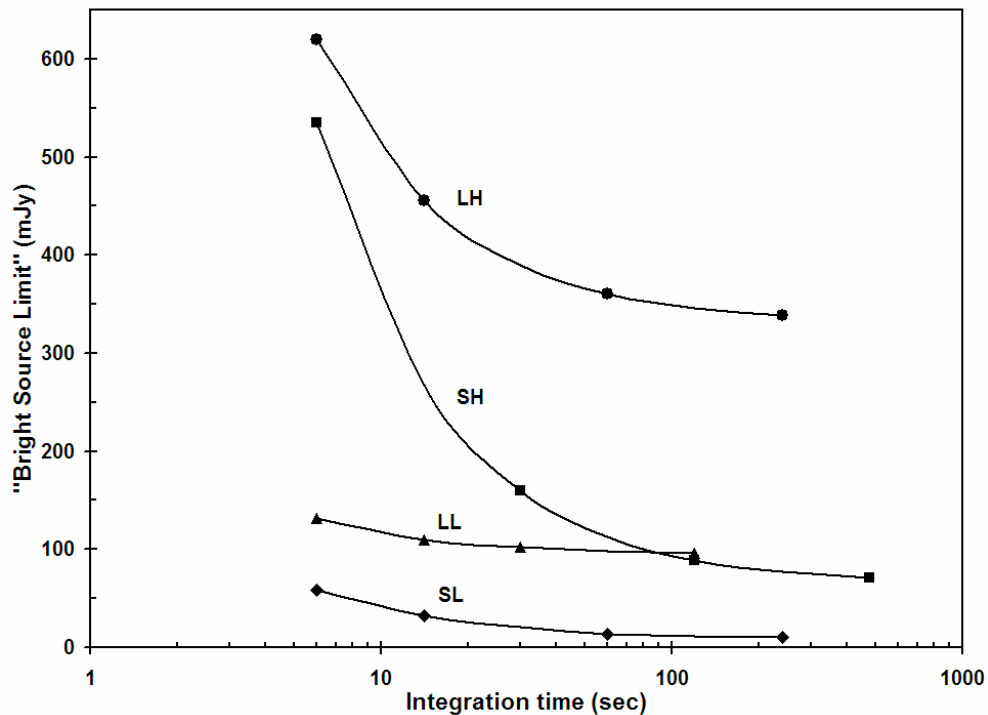
Figure 7.24: Bright Source Limit (BSL) as a function of integration time for each of the IRS modules. Targets with fluxes above the curves for a given integration time are source-dominated; targets with fluxes below the curves are read noise or background limited. The plotted points on the curve for each module indicate the allowed AOT durations for that module.

7.2.1.2.2   *Estimating S/N for Bright Sources*

It is important to realize that the PSSC and PSSCS as calculated in section 7.2.1.1 and shown in the sensitivity curves in section 7.2.1.2 assume that the target is in the faint source limit; that is, the shot noise from the source is negligible compared with the other sources of noise described above. Conversely, Figure 7.24 shows the approximate flux (the median flux averaged over the whole wavelength response of each module) at which a target's shot noise dominates over the other sources of noise (sky background + dark current + read noise), as a function of exposure duration for each of the IRS modules. Thus, Figure 7.24 can

be used to determine whether a target object of known IR flux corresponds more appropriately to the "Bright Source Limit" (BSL).

The low resolution modules rapidly reach an asymptotic BSL in a few tens of seconds of integration time, whereas the high resolution modules take longer to reach this plateau. Sources that have fluxes which lie below the curves plotted in Figure 7.24 for a given integration time are in the faint source limit (i.e., they are background or read noise limited), whereas sources that have fluxes which place them above the line for a given module have signal-to-noise ratios (S/N) that are dominated by the source flux. Knowledge of the regime in which a source lies is crucial to calculating a realistic S/N estimate needed to plan an IRS observation. Note that for all allowed AOT integration times of the low resolution modules, sources brighter than 130 mJy will always be source-dominated. For all allowed AOT integration times of the high resolution modules, sources brighter than 620 mJy will always be source-dominated.

In this section, we provide a simple scaling rule that allows the original estimate of the signal-to-noise ratio for faint sources obtained through the use of Figure 7.19–Figure 7.22 (see section 7.2.1.2) to be corrected for the contribution of the shot noise from a bright source. Consider, for example, a continuum source with a known 8 μm flux of 250 mJy. To estimate the S/N of this source in the SL module for an integration time of 14 sec, we first consider the PSSC for SL at this wavelength (i.e., the faint source limit case). Inspection of Figure 7.19 shows that at 8 μm, the 1σ PSSCS is 0.7 mJy (in the LOW background case). If we were to (incorrectly) assume that the source itself contributed no noise, then we would estimate the signal-to-noise ratio $(S/N)_{faint} = (250/0.7) = 357$. However, we notice from Figure 7.24 that in SL for a 14 sec integration time, the Bright Source Limit (BSL) = 26.5 mJy. Thus, the 250 mJy target exceeds the BSL by a factor of ~9.4. We can now estimate a more realistic S/N for the source by using the simple scaling rule

$$(S/N)_{bright} = (S/N)_{faint} /(1+[\text{Source Flux/BSL}])^{0.5}, \tag{7.11}$$

where the source and BSL fluxes are given in the same units (e.g., mJy). In this example,

$$(S/N)_{bright} = 357/(10.4)^{0.5} = 357/3.2 = 112. \tag{7.12}$$

Hence, a more realistic value for the 8 μm S/N for a 250 mJy source observed for 14 sec in SL is ~110. This procedure can be generalized for use in estimating the S/N that will be obtained for observations of bright sources (both point and extended — see section 7.2.1.2.1) using any of the IRS modules.
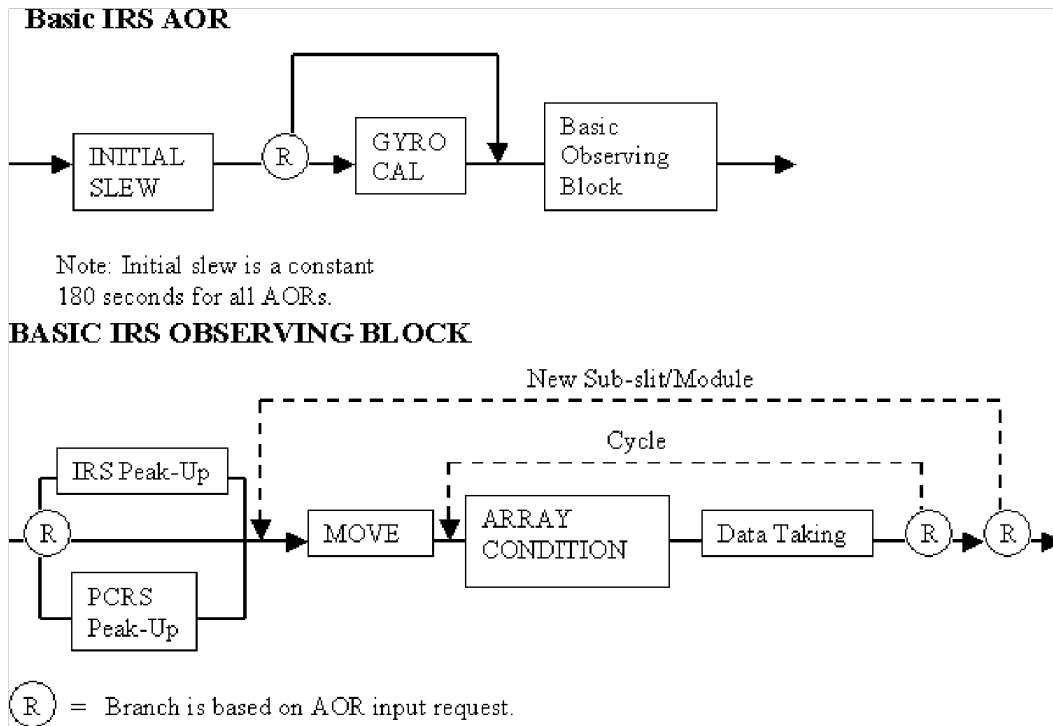
Figure 7.25: Block diagrams showing a generic timeline for an IRS AOR.
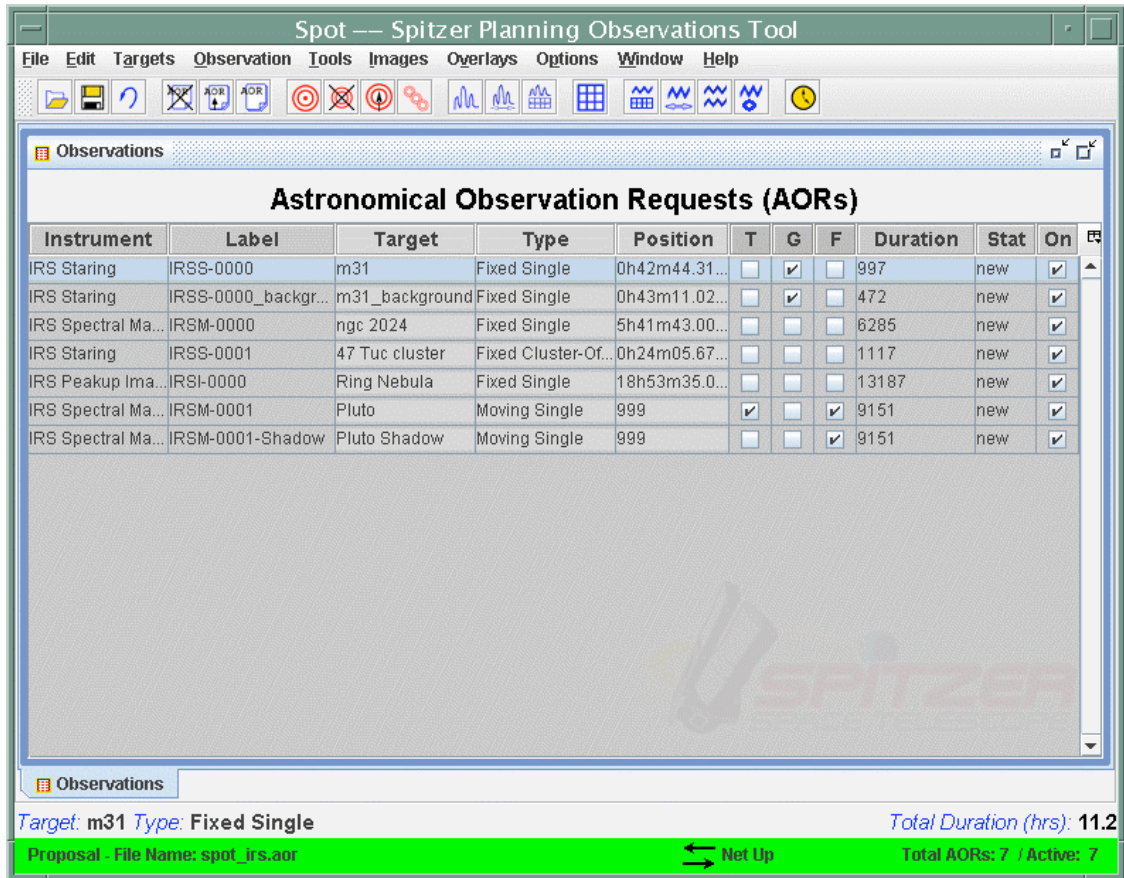
7.2.2 Observing Time Estimation

There are two basic time components associated with IRS observing. The first is the integration time, or the time "on source". This is the time spent collecting useful photons, and does not include such activities as peak-up or moving the target between adjacent slit positions. The second time component is observational overhead. This includes the reset and conditioning frames taken at the start of each ramp, the peak-up, all telescope motions within the AOR, and a generic slew overhead to account for the time it takes for the initial (large) slew to the peak-up or science target. Spot calculates all of the necessary overhead times after the observer selects the desired integration times. Together, the integration and overhead times are added to estimate the total AOR duration, which is the amount of time that an observer requests in his/her proposal. Figure 7.25 shows the timeline for carrying out an IRS AOR in a block diagram format. Observers should note that the array conditioning time may be "hidden." That is, it may begin, and be completed, during the move.

7.2.3 Astronomical Observation Template Description

Astronomical Observation Templates (AOTs) are the building blocks that observers use to plan and request Spitzer observations. Each of the three Spitzer instruments has a number of available AOTs, and each of these AOTs is accessed through Spot, from either the Observation menu or the specialized AOT icons on the main Spot window (see Figure 7.26). For the IRS, the available AOTs are: Staring, Spectral Mapping, and Peak-Up Imaging; these are described in sections 7.2.3.2, 7.2.3.3, and 7.2.3.4, respectively. In general, an AOT consists of a form

that specifies information about the target, observing mode, and instrument set-up parameters. A completed and validated AOT results in an Astronomical Observation Request (AOR), which is essentially a Spitzer command sequence designed to perform the observation. Because of required instrument calibration events, **an IRS AOR can have a total duration (including all overheads) of no longer than 8 hours.** In the following sections, we describe the parts of the IRS AOTs. Some parts are common to all IRS AOTs, while others are found only in specific AOTs. (Peak-up target acquisition for the Staring and Spectral Mapping AOTs is a potentially complex issue and is discussed separately in section 7.2.4.)

Figure 7.26: The main Spot interface window, showing a list of AORs that summarizes the information needed to perform a sequence of observations.

7.2.3.1 Target Specification

The first entry in all of the IRS AOT forms in Spot (and one that is common to all Spitzer AOTs) is the information needed to specify the science target. Pressing the New Target button at the top of the AOT form generates a window (see Figure 7.27) that allows entry of a target name and selection of one of five target types: (1) Fixed Single, (2) Fixed Cluster-offsets, (3) Fixed Cluster-positions, (4) Moving Single, or (5) Moving Cluster. Fixed Single is the "basic" IRS target type, corresponding to a single, stationary object, specified by a set of celestial coordinates, coordinate epoch, and proper motions. Moving Single is the

analogue for non-stationary (i.e., Solar System) targets, which are specified by their NAIF number and orbital ephemeris. The Fixed Cluster target types allow multiple, identical IRS observations of several targets (or several positions in an extended target) within a radius of less than 1 degree from a central target. This target type has advantages in reduced overhead compared to observing each position using separate Fixed Single targets. The suffix "offsets" or "positions" denotes the format used to specify the target positions for the Fixed Cluster target types, either as offsets from a central target position or as a set of celestial coordinates for each target position, respectively. The Moving Cluster target type is analogous to the Fixed Cluster-offsets target type, but applies to non-stationary targets.

Figure 7.27: Target entry window.

The visibility of a target to Spitzer can be checked by clicking on the "Visibility/Orientation" button in the Target panel. There are two bright Solar System target avoidance boxes in the Target panel. One is labeled "Earth + Moon", and (obviously) includes the Earth and the Moon. The second avoidance box is labeled "Others" and includes the planets (except Pluto) and the brightest asteroids visible to Spitzer. Checking these boxes ensures that a target will only be observed when it is more than 7 degrees (Earth + Moon) and/or 30 arcminutes (Others) away from any of the excluded objects – see the Spot User's Guide for a more complete discussion.

Once the target information is entered, close the target window by pressing the OK button. After all target(s) have been entered into the AOT, a Target List is generated by Spot. All targets in this list are accessible to the observer for any of the Spitzer instruments and AOTs, and the observer can select targets from this list for one or more observations. If targets are added later in the same Spot session during configuration of additional AOTs, then these are appended to the existing target list. Targets may also be entered prior to selecting an AOT by selecting the Target icon from the main Spot window, or by reading in a pre-assembled list of targets (e.g., saved from the Target List in a previous Spot

session).

7.2.3.2 Staring Mode

Staring mode is the basic "point and shoot" operating mode of the IRS. In this mode, science targets are placed on one or more of the IRS slits for a specified integration time. The target parameters, desired spectral wavelengths (slits), and method of target acquisition (i.e., peak-up — see section 7.2.4) are input to the AOT form (see Figure 7.28) by the observer. An estimate of the total observation time, the overhead times, the observation visibility windows and slit position angles (PAs) for each window, and a means to save the planned observation in ASCII format are all provided by the Spot software.

7.2.3.2.1   Instrument Settings for Staring Mode

The IRS Staring mode AOT can be configured as either Standard or Peak-Up Only by selecting the appropriate tab in the Instrument Settings section.

7.2.3.2.1.1   Standard Staring Mode

In the Instrument Settings section of the Standard Staring mode AOT, the observer specifies the spectra to be obtained with the IRS. Each IRS slit corresponds to a particular wavelength region and a particular spectral resolution. There are six "slits" (or "modules") to choose from since SL and LL have two sub-slits each. The observer must enable a slit for an AOT by turning it "on" in Spot. Note that Spot uses the formal names for the modules; see Table 7.2 for a list of names used for each module. For each slit, the observer must select from one of four ramp durations, and then enter the number of cycles of that ramp duration to be performed. The ramp duration is the time between the first and last non-destructive reads of the array, and corresponds to the "effective exposure time" in Figure 7.4. The shortest ramp duration for all of the slits is six seconds. The longest ramp duration for the slits varies from 120 seconds to 480 seconds depending upon the module. The ramp durations have been chosen by the IRS Instrument Team to provide the highest dynamic range, while minimizing the number of routine calibrations required in orbit. Note that the IRS SL and LL modules reach background-limited performance in about 40 seconds and about 10 seconds, respectively. Longer ramp durations have smaller overhead times but may be affected more severely by cosmic ray hits. The "cycles" parameter is the number of times a given spectrum is repeated before moving on to the next slit position or the next target. Total "integration time" is determined by both ramp duration and number of cycles.

The automatic "nodding" that occurs in Standard Staring mode means that, for example, the choice of a 480 second ramp for one cycle in SH would produce two spectra, with each having an exposure time of 480 seconds, yielding a total on-source integration time of 960 seconds. The same exposure time with two cycles would produce four 480-second exposures, for a total on-source integration time of 1920 seconds.

Figure 7.28: Staring mode AOT entry form.

Standard Staring mode successively places the target at the two nominal slit locations for each requested slit, located 1/3 and 2/3 of the way along the length of the slit, to provide redundancy against cosmic rays and detector artifacts.

7.2.3.2.1.2   Peak-Up Only Staring Mode

The Peak-Up Only observation type is intended to provide "early target acquisition" functionality to the IRS Standard Staring mode. In a Peak-Up Only observation, the peak-up algorithm is run in its normal fashion to acquire the science target (see section 7.2.4.1.6) but no spectroscopic data are taken. This may be useful in cases where long spectroscopic observations are needed but no obvious peak-up target presents itself from other ground-based or space-based

data, so it is desirable to confirm the success of the peak-up before exposing the spectrum. The peak-up images and the position of the target selected by the software are sent to the ground in this mode. Since the peak-up software runs its normal course here, the overheads (including the initial slew overhead and the settling times within the peak-up process) are significant. In addition, the peak-up itself will have to be repeated at a later time in conjunction with the actual spectroscopic exposure. The DCS images obtained during a Peak-Up Only observation will be dark-current-subtracted, cleaned of cosmic rays, and flat-fielded on-board (as is done for all peak-up data) and these images will be made available to the observer for planning future IRS spectroscopic observations. Peak-Up Only is intended to test the peak-up algorithm, not to obtain science images, which can be accomplished with the Peak-Up Imaging AOT (see section 7.2.3.4).

Figure 7.29: Spectral Mapping mode AOT entry form.

7.2.3.3 Spectral Mapping Mode

The Spectral Mapping AOT (see Figure 7.29) allows an observer to configure a grid of map positions around a central target position and obtain spectra at each position. The observer may specify any combination of slits, each with different ramp durations and cycles, for single or cluster targets (provided the 8-hour AOR maximum is not exceeded). The telescope is moved in a direction perpendicular to the long axis of the slit for the desired number of steps before moving parallel to the slit for the desired number of steps. By default, each map position produces a single ramp (unlike Staring mode, which automatically obtains spectra at two nod positions along the slit), but the observer can set a larger number of ramp cycles to be obtained at each map position or repeat the entire mapping sequence multiple times. The individual pointings are arranged so that the science target defines the center of the map. The orientation of the map on the sky is determined by the date and time of observation.

7.2.3.3.1   Instrument Settings for Spectral Mapping Mode

The observer can choose any combination of IRS slits in Spectral Mapping mode. The general information regarding slit selection given in section 7.2.3.2.1.1 for Standard Staring mode also applies for Spectral Mapping mode. For example, if you click on the "Low 5.2–8.7" and "Low 7.4–14.5" buttons, then you will get two maps, one centered on the 2nd order SL slit and one centered on the 1st order SL slit. In this case there will be many extraneous spectra taken in both orders since there is only one physical SL slit (so spectra are always obtained in both orders – with one order offset from the target coordinates – even when only the 1st or 2nd order subslit is selected). In addition, for efficient mapping of extended sources, there are two more module choices in Spectral Mapping mode: Low Short Both and Low Long Both. These use the vignetted region between the subslits of the indicated modules as the reference point, rather than the center position along each subslit. If maps are desired in both 1st and 2nd orders, this option offers a significant savings in time over mapping in each subslit separately (yet see below for possible dangers in using the "Both" option). Regardless of which module is selected, the observer must specify the step size in arcseconds and the number of steps for each leg of the map. When observing a cluster with more than one slit in Spectral Mapping mode, all targets are observed first in one slit, then all targets are observed in the second slit, and so on. This yields better control of the pointing accuracy while observing in the narrow slits, and reduces the overhead necessary to switch from one IRS module to the next.

The observer can choose to map a given target (or cluster of targets) more than once using the "Number of Map Cycles" parameter (see Figure 7.29). This will invoke a peak-up (if requested), reposition the telescope to the beginning of the first map leg of the first target, and repeat the entire map, starting with the narrowest slit requested. Mapping large areas will typically not require high pointing accuracy, but observers are reminded that the telescope pointing suffers random positional offsets of order 0.2″ for every commanded movement of the telescope. Spectral mapping of multiple sources with multiple slits can generate a

very large number of telescope movements, with significant associated AOR duration overheads if high positional accuracy is requested.
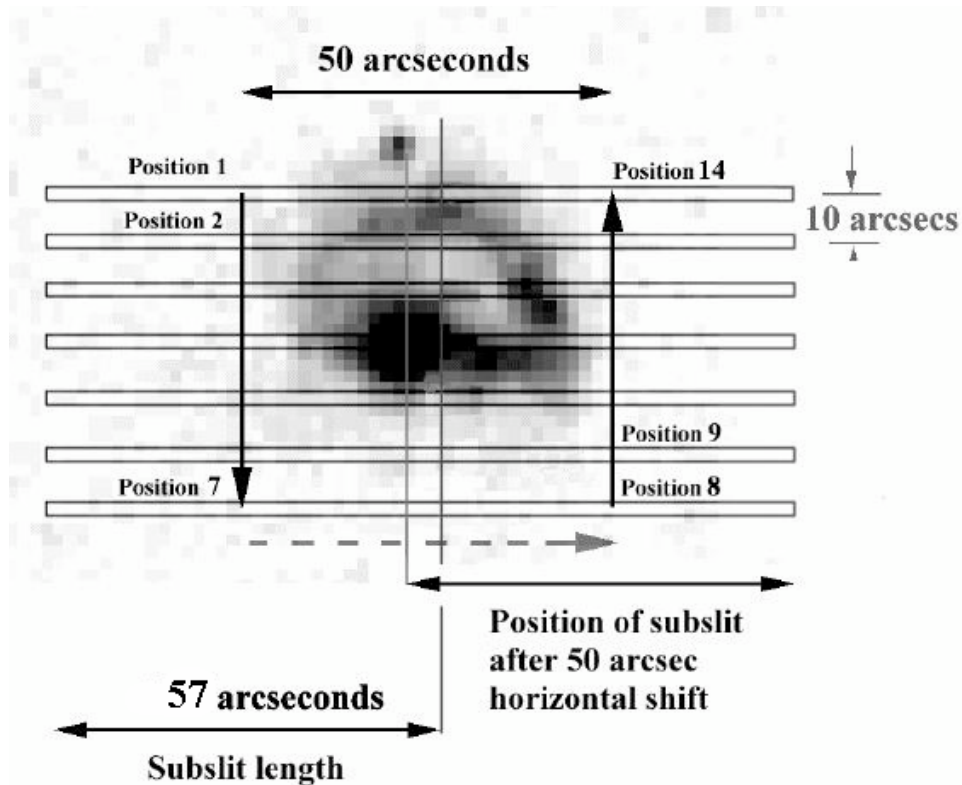
Figure 7.30: Example of a sparse 7×2 Spectral Mapping grid that ensures that the slit passes through the nucleus of the galaxy. Evenly numbered (e.g., 8×2) sparse grids will not place a slit at the center of the map (which is always centered at the target coordinates — in this case, the galaxy nucleus). This example demonstrates an additional danger; namely, that the grid is so sparse it does not efficiently sample the area of interest.

7.2.3.3.2   *Potential Pitfalls in the Use of Spectral Mapping*

Given the freedom to specify arbitrary step sizes (and, therefore, potentially "sparse" or unfilled maps), care should be taken to ensure that the field of interest is mapped as intended. Hence, we caution the observer regarding the following facts about Spectral Mapping.

If a sparse map is configured (i.e., if the steps are larger than the slit width) and the observer wants to have the slit pass through the position of the map-center (for example, centered on the nucleus of a galaxy), then an **odd number of steps** in the "perpendicular" direction is necessary; otherwise, the central source will fall between two slit positions. Figure 7.30 shows an example in which 2 parallel and 7 perpendicular observations are made in Spectral Mapping mode (7×2 grid). Similarly, when using either the "Low Short Both" or "Low Long Both" options, specifying a large slit-parallel step size and an with an odd number of steps could put the center of the map (i.e., the central target coordinates) in the vignetted region between the subslits.

It is almost always useful to overlay the final mapping pattern on the field of interest using the visualization tool animation sequence in Spot. See section 7.2.7.3 for an IRS Spectral Mapping mode example and the Spitzer Observation Planning Cookbook for additional discussion about designing a successful Spectral Mapping AOR. A primer on the design of dense IRS spectral maps is also given at http://ssc.spitzer.caltech.edu/irs/documents/specmap_bop/.

7.2.3.4 Peak-Up Imaging Mode

The IRS Peak-Up arrays provide a science-quality imaging capability in the wavelength range that falls between that covered by IRAC and MIPS, and as such may be of interest for imaging alone. Starting in Cycle-2, a new AOT was made available to allow Peak-Up Imaging (PUI) observations. Both a red (18.5–26.0 µm) and a blue (13.3–18.7 µm) filter are available (see Figure 7.38). The red filter provides images over a similar wavelength window as does the MIPS 24 µm mode. The main motivation for the AOT is to provide photometric imaging in the 16 µm (blue) filter, thereby filling the wavelength gap between IRAC and MIPS. Unlike the peak-up images taken for the purpose of slit centering, the peak-up imaging mode produces SUR data.

Point source saturation limits are provided in Table 7.7. Sensitivities are provided in Table 7.6.

7.2.3.4.1   Peak-Up Image Characteristics

Parallel red and blue peak-up images are obtained simultaneously. The field-of-view seen by each filter is 54″×81″, and they are separated by a 33″ wide vignetted zone (see Figure 7.6). The images are 35×45 pixels in size and have a plate scale of ~1.8″/pixel. The position of the two images on the SL detector is an accurate representation of their position on the sky; that is, the two fields-of-view together with the vignetted region see a contiguous 54″×195″ area. A single World Coordinate System (WCS) can be used to describe the combined red and blue regions within the 128×128 pixel detector.
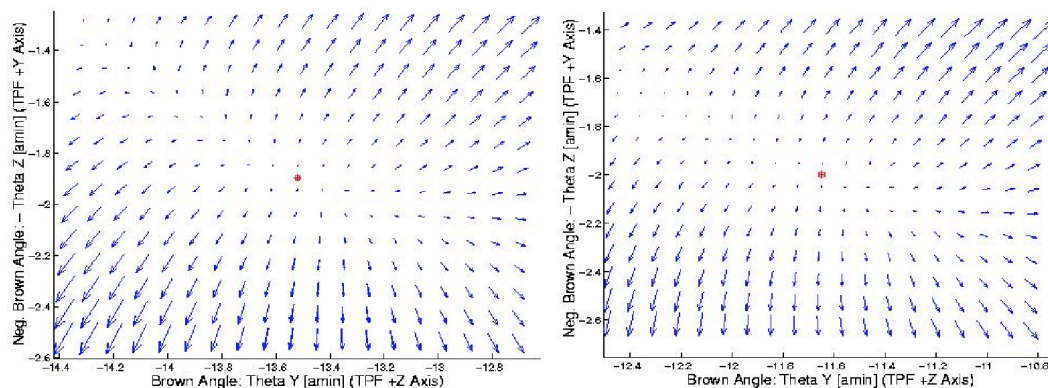
Figure 7.31: The relative magnitude and direction of the geometric distortion in the blue (left) and red (right) IRS Peak-Up arrays. The distortion is plotted at 5× magnification.

7.2.3.4.2   *Peak-Up Image Quality*

The IRS provides diffraction-limited images through both peak-up filters. The geometric distortion has been accurately mapped and found to be small (see Figure 7.31). The plate scale is 1.85″/pixel in the detector X direction, and 1.82″/pixel in the detector Y direction. The magnitude of the distortion causes no more than a 1/2 pixel shift at the edge compared to the center of either peak-up array field-of-view. The distortion coefficients are applied to the astrometric solution in the WCS. **The full-width-at-half-maximum (FWHM) of the point spread function (PSF) is ~2 pixels (3.8″) in the blue filter, and ~3 pixels (5.3″) in the red filter.** The radius containing 50% of the encircled energy is 1.6 pixels (3.0″) in the blue filter, and 2.2 pixels (3.9″) in the red filter. These numbers were derived for stellar (blue) sources. Figure 7.32 shows images of the IRS Peak-Up PSFs.
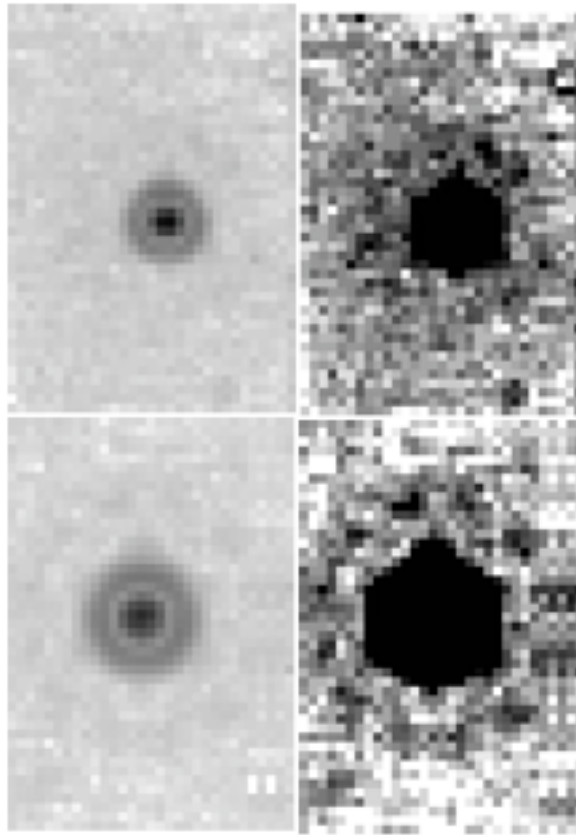
Figure 7.32: PUI PSFs for the blue (top) and red (bottom) apertures. These PSFs are co-additions of several high signal-to-noise stellar images selected to lie within 0.1 pixels of one another. The images on the left are stretched to show the central PSF and the first diffraction ring; the images on the right are stretched to show the second diffraction ring. These images show the full fields of view of the peak-up arrays.

7.2.3.4.3   *Parallel Spectra*

The peak-up apertures are imaged onto the IRS SL detector array. A read-out of the array will therefore, also yield SL1 and SL2 spectra at parallel locations on the sky. The locations of the slits are determined by the orientation of the telescope,

which depends on the date of observation. Nonetheless, science quality spectra are obtained. This mode may be particularly useful for sparce spectral mapping or looking for serendipitous detections of rare objects. The similarity in PUI and spectroscopy data reduction procedures ensures that the observer will receive fully calibrated 2D spectra along with the imaging data.

7.2.3.4.4   Peak-Up Imaging Sensitivity

We use the same basic calculations presented in section 7.2.1.1 to estimate the sensitivity and S/N for PUI observations. Table 7.6 summarizes the instrument sensitivity in the PUI mode and can be used for planning PUI observations. The flux density is the peak-up target flux density at 16 μm (blue) or 22 μm (red).

Table 7.6: PUI 1σ sensitivity (μJy) for 6/14/30 second ramp times

Filter	LOW background	MEDIUM background	HIGH background
Blue	115 / 75 / 50	145 / 95 / 65	190 / 120 / 80
Red	180 / 115 / 75	220 / 140 / 95	280 / 180 / 120

The underlying procedure used to calculate the PUI sensitivities is presented for informative purposes here:

1. The background photocurrent from sky and telescope is given by Equation 7.1, where, in the case of PUI, $\Delta\lambda \approx 5.4$ μm.

2. For faint sources, the total noise in electrons per pixel is the sum of the photocurrent shot noise, read noise, and dark current shot noise; it is given by:

$$N = \sqrt{i_{sky}t_{int} + (RN)^2 / N_{reads} + i_d t_{int})} \quad [\text{e- px}^{-1}] \tag{7.13}$$

where i_d is the dark current contribution, RN is the read noise in electrons, N_{reads} is the number of nondestructive samples up the ramp, and t_{int} is the ramp time in seconds.

3. The relationship between photocurrent per pixel and the incident flux density, S_{mJy}, is given by equation 7.4 except that the slit losses, τ_{slit}, are replaced by the fraction of encircled energy in the assumed beam size, which, for PUI, is larger than the 4 pixels assumed in Equation 7.4. For the nominal S/N calculations, we assumed a beam size equal to that of a circular aperture containing half the light of a point source.

4. The S/N for a point source is then the ratio of S_{mJy} to N.

7.2.3.4.5   Peak-Up Imaging Data

As shown in Figure 7.12, peak-up images are read out on the SL detector. In the PUI mode, the SUR data collection mode is employed (see section 7.1.3.1.2). SUR observations are preferable to the DCS observations used by the IRS peak-up target acquisition mode because the former substantially reduces read noise

and, more importantly, allows for cosmic ray rejection in a single image.

7.2.3.4.5.1   Data Processing

The basic data processing steps for PUI are virtually identical to those for spectroscopy on the SL detector. These steps are described elsewhere in this chapter. Only after application of the flat field are imaging-specific data reduction steps required.

Separate calibration observations are performed to obtain the relevant data (e.g., flat field images), which are then combined into composite calibration FITS files.

Sky Darks: As with spectroscopic observations, multiple observations of low background regions of the sky are combined into a calibration file to be subtracted, plane-by-plane in the SUR data, from the science observations. While this "sky dark" calibration is meant primarily to remove the detector bias voltage and the (relatively small) dark current that accumulates during the exposure, sky emission is also subtracted. However, the zodiacal background is a strong function of position on the sky, and a weak function of the time when the data were taken. As a result, many calibrated observations will still have substantial residual sky background. As with all infrared data, it is necessary to plan the observations to provide sufficient measurement of the local sky at the time. Note that the variation of the zodiacal emission during the year may also result in observations in which the standard calibration will over-subtract the background.

Sky Flats: The flat fields are generated from observations of regions of extended emission (e.g., high zodiacal light background). The flat fields are very uniform over the peak-up array fields-of-view, with a standard deviation in relative pixel response of ~1.2%.

Photometric Calibration: The flux calibration of the PUI science mode is updated periodically. Measurements of a set of standard stars are compared to stellar models to derive transformations from electrons/sec to Jy.

Image Cropping and Flux Calibration: The final stage of the science data processing pipeline for PUI is to crop the red and blue Peak-Up images into individual 35×45 pixel FITS files. Each file has the flux calibration applied to put it in units of MJy/sr, and an individual WCS in the header.

Mosaic Registration: The data reduction pipeline at the SSC does not, at this time, provide registered mosaics of dithered observations. However, the individual processed data products are appropriate for use with the MOPEX post-BCD software (see the SSC website).

7.2.3.4.5.2   Data Products

Additional details regarding the PUI data products are available in the IRS Data Handbook. A brief summary is presented here.

Raw Data: The raw data are provided, without calibrations applied, as 3D, 16-bit integer data cubes in FITS format. Each plane of the cube corresponds to one of the nondestructive reads in the sample-up-the-ramp. The FITS header contains the extended telemetry information, but the values of use to the general observer are given "readable" keyword names in the processed data header.

2D Processed Data: The full 128×128 processed 2D data frame is provided. This FITS image is a single plane, in units of electrons per second, and contains the red and blue PUI data, together with the parallel SL spectra. Two associated FITS files are also provided: the uncertainty plane and the bmask (the bit-flag pixel status mask). A single WCS is used in the 2D image, based on the blue peak-up image astrometric solution. The geometric distortion will, therefore, not be as accurate for the red field-of-view.

Cropped FITS images: The red and blue peak-up data are cropped into "postage stamp" FITS files (see section 7.2.3.4.5.1), with individual WCSs. Each file has associated uncertainty and bmask files.

7.2.3.4.5.3   Latent Images

Bright objects falling on the IRS SL detector can result in persistent charge that appears as a latent image of the object in subsequent exposures. For exposures with $10^4-5\times10^4$ e- pix^{-1} in the last sample, 0.04 to 0.1% will remain behind as latent images in the subsequent exposure. These latents decay away rather slowly, with an e-folding time of ~1100 sec, and for many successive exposures source flux can appear to increase by several percent. Given the parallel position of the peak-up fields-of-view to the SL slits, it is not uncommon for a bright object to fall on the peak-up arrays during science operations. The observer should take this effect into account in planning observations within each AOR. There is no way to avoid the latent images caused by previous science programs; however, IRS observing and data-acquisition strategies are designed to minimize this aspect of the latency effect. Proper dithering strategies will mitigate the latent image problem, although with some loss of signal-to-noise on the affected pixels.

7.2.3.4.6   How to Use Peak-Up Imaging

The PUI AOT in Spot is designed to be familiar to observers who have previously configured IRAC and/or IRS Spectral Mapping AOTs. Observations consist of an optional dither pattern superposed on an optional rectangular map pattern. The exposure time and dither/map patterns are defined in a single dialog box in Spot (see Figure 7.33). When mapping and dithering are selected, the required dialog fields open within the main box. PUI observations may be requested for any of the standard target types (see section 7.2.3.4.6.3 for an exception to this for PUI Mapping). An example PUI observation is described in section 7.2.7.5.

7.2.3.4.6.1   Ramp Times and Cycles

There are three selectable ramp (exposure) times for PUI: 6, 14, and 30 seconds, corresponding to 4, 8, and 16 nondestructive reads in SUR mode. As in

spectroscopy, more samples allow for better cosmic ray rejection and reduce read noise (although PUI observations are heavily background limited). It is important to consider saturation when selecting a ramp time, especially in cases of bright source flux or high background. Table 7.7 shows the point source saturation limits for the PUI mode, depending on the source brightness at 16 μm (blue) or 22 μm (red). The extended source saturation limit at 6/14/30-second ramp times occurs at 450/190/90 MJy/sr in both fields. The fluxes listed in the table will saturate the detector at the top of the ramp, but data processing can recover the slope of the ramp (at lower S/N) as long as the detector does not saturate the 4th (or even 3rd) read. Multiple ramps can be selected by choosing a number of cycles > 1. Data are collected for all cycles before offsetting the telescope. Each cycle is an independent DCE, and will increment the DCE number in the name of the FITS files.

Figure 7.33: PUI AOT form in Spot, without dithering or mapping.

Table 7.7: Point source saturation limits (mJy) for 6/14/30 second ramp times

Filter	LOW background	MEDIUM background	HIGH background
Blue	190 / 90 / 45	190 / 90 / 40	180 / 80 / 35
Red	430 / 210 / 100	420 / 200 / 90	410 / 190 / 80

7.2.3.4.6.2   Fields-of-View

PUI observations may be requested in either red or blue or both. If both red and blue are selected, then all observations (including dithering and mapping) will be performed in sequence, first with the blue aperture and then with the red. Each peak-up field-of-view can have its own integration parameters (ramp times and cycles). It is also possible to select the "middle" field-of-view, which will use the center of the vignetted region between the blue and red apertures as the targeting reference. If the middle field-of-view is selected, then only a single set of dither/map positions is used and the integration parameters specified for the blue aperture will apply to the red, as well. It is not possible to select either red or blue in addition to middle. The use of the middle option can make the mapping of extended areas more efficient, taking half as much time and generating half as much data as separate blue and red maps, and is included for consistency with IRAC. If neither dithering nor mapping is selected, then the observation will consist of single pointings in each requested field-of-view.

7.2.3.4.6.3   Mapping

The mapping grid consists of a rectangle of positions, specified in terms of rows and columns with an observer-selected spacing (see Figure 7.34). The grid is centered on the target position. An odd number of map positions in rows and/or columns will result in an observation centered on the target in the row and/or column direction; conversely, an even number of map positions results in the target being centered *between* map positions. If desired, the entire map can be repeated as many times as the 8-hour IRS AOR limit will allow. **Note that this is neither a preferred nor an efficient means of increasing integration time.** Specifying multiple exposures per map position using the "# of cycles" parameter will yield any desired total integration time, and reduce by a factor of "#", the number of telescope movements required.

PUI map positions may be specified in either array coordinates or celestial coordinates. If array coordinates are selected, then rows are aligned with the X axis of the detector and columns with the Y axis. In celestial coordinates, the overall position angle of the map can be selected. The celestial coordinates option is particularly useful in mapping regions that are highly asymmetric on the sky. It also has the advantage of fully specifying grid locations in multiple AORs independent of observing date. However, celestial coordinate grids can be subject to significant coverage gaps because the rotation of the individual peak-up apertures on the sky still depends on the date of observation. PUI maps can be visualized in Spot, but, in most cases, the scheduling of an observation (and, hence, the orientation of the peak-up apertures) is not known when the AOR is designed. However, the AOR overlay can be examined for the full range of the target's visibility window, allowing maps to be optimized for all possible orientations.

An offset in each of the grid directions can be specified. This option is useful for very large maps. IRS AORs are limited to durations of 8 hours, so longer maps must be broken into smaller pieces. Using the map center offset option makes it

easier to ensure the desired spacing between AORs. Care must still be taken to account for orientation angles that change with observing date. The offset option is also useful to obtain local sky observations by offsetting from an extended target.

> **Note that PUI mapping is *not* available for target types of Fixed Cluster-offsets or Fixed Cluster-positions.**

Figure 7.34: PUI AOT form in Spot, with mapping selected.

7.2.3.4.6.4   Dither Patterns

Dithering is designed to allow the sampling of a given sky position by multiple pixels. This helps to remove detector problems such as bad pixels, latent images, or pixel-to-pixel errors in the flat field. Well-dithered observations can also be used to measure the sky background in fields without extended sources. Finally, sub-pixel dithering can be used to recover additional spatial information (though the improvement is larger in under-sampled data, rather than the critically-sampled blue peak-up array).

Two sets of standard dither patterns for PUI are available: random and cycling (see Figure 7.35 and Figure 7.36). The random patterns consist of five or nine

positions generated from uniform random distributions. Large and small steps are available. The small pattern has a minimum step size of 2 pixels, which is the FWHM of point sources in the blue peak-up array. The area of full overlap in registered mosaics using the small patterns is about twice that of the large patterns. The cycling patterns consist of a table of 311 offsets that can be used starting at any point in the list. They are drawn from a Gaussian distribution, but restricted so that no group of four positions is closer than 2 pixels to each other. The cycling patterns are designed to allow the selection of an arbitrary number of dithers. They are particularly useful in combination with mapping, to allow a small offset from the regular spacing, and improve the sampling of the flat field and any potential pattern noise.

Nine point patterns have dither positions with 1/3 sub-pixel sampling. That is, offsets are performed in multiples of 1/3 of a pixel. Five point patterns have 1/2-pixel dithers, but an additional 1/4-pixel offset is added to the fifth point. The cycling patterns have 1/2-pixel steps.

Figure 7.35: PUI AOT form in Spot, with dithering and mapping selected.
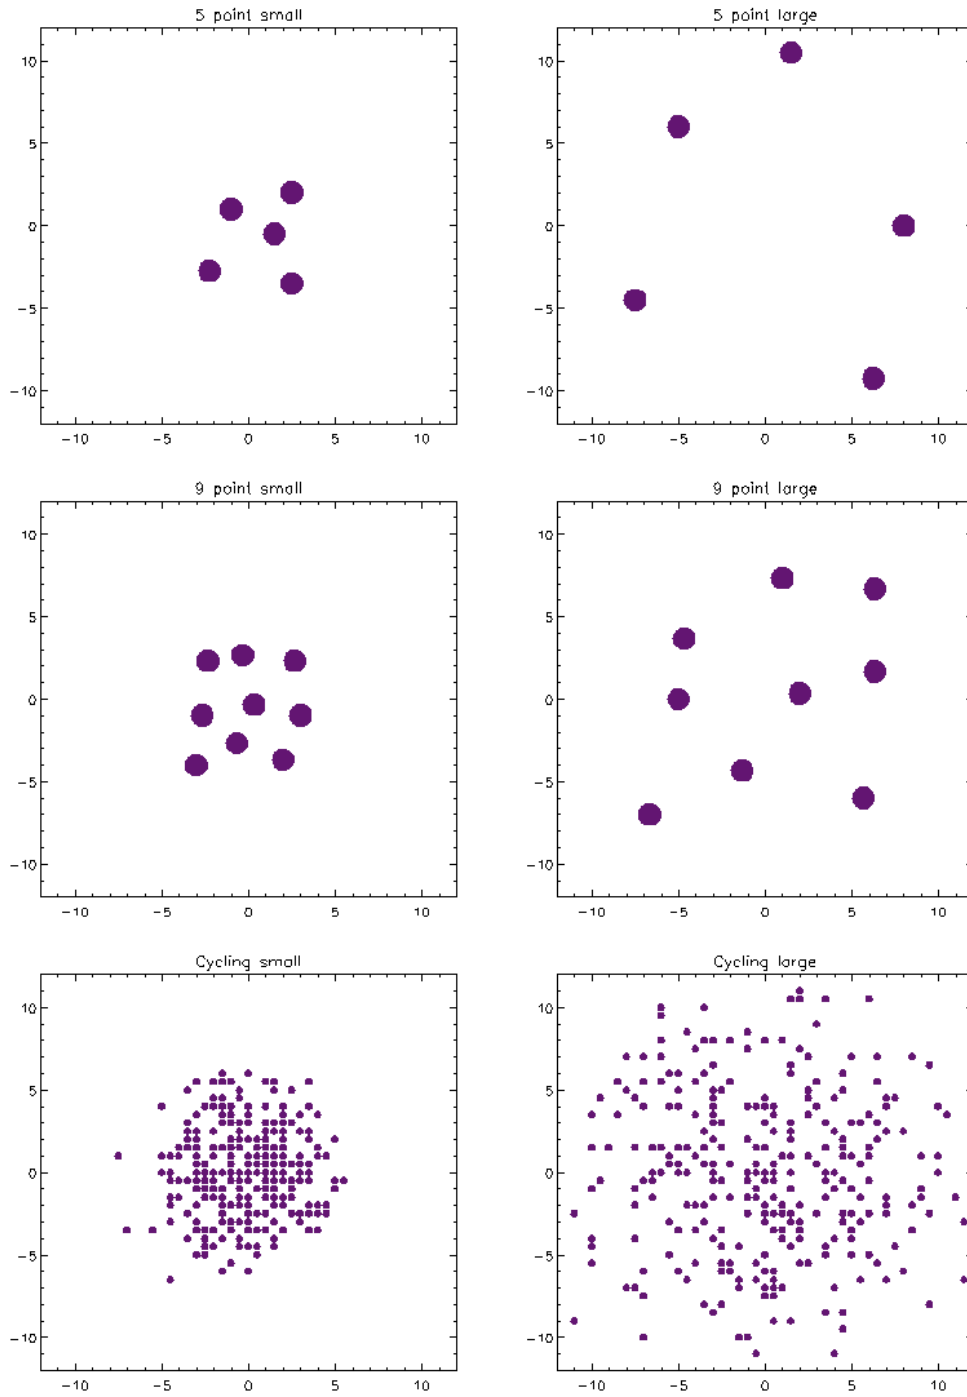
Figure 7.36: The available PUI dither patterns.

7.2.3.5 Calculating Observing Time

The Calc. Obs. Time button uses the target and instrument specifications to compute the following quantities: (1) slew time in seconds, (2) settle time in seconds, (3) total integration time in each slit in seconds, and (4) total AOR duration in seconds (see Figure 7.37). The total duration includes all overheads within the AOR, including a generic slew overhead of 215 seconds that accounts

for the average time required to slew to the region of the sky specified in the AOR (an additional 300 sec overhead is charged for moving targets). These overheads may include some or all of the frame conditioning times associated with SUR data collection (if frame conditioning can not be carried out while slewing and settling), the time taken for any intermediate peak-ups required to maintain the requested level of pointing accuracy, inertial hold times necessary to achieve and maintain the highest pointing accuracy, as well as smaller overheads of typically a few seconds for initial setup, switching between modules, and so forth.

> **Note that requested ramp times and actual integration times differ slightly because Spot takes into account the difference between the IRS clock rate used to measure the ramp times and the main spacecraft clock, which measures the actual integration time in standard seconds. The integration times reported by Spot are the ones to use in estimating total exposure time.**

Figure 7.37: Sample IRS resource estimation window.

> **Special Overhead is an additional charge that reflects the scheduling cost of special operations such as Targets of Opportunity, three-instrument campaigns, or late ephemeris updates for moving targets.**

7.2.4 Peak-Up Target Acquisition

An important step in the preparation of Staring mode and Spectral Mapping mode AORs is consideration of the peak-up target acquisition settings, which are found in the middle of the AOT entry forms (see Figure 7.28 and Figure 7.29). Peak-Up Imaging mode does **not** use a peak-up target acquisition. The motivation for having an on-board peak-up capability is essentially two-fold. First, since the IRS slits are relatively narrow compared to the nominal pointing accuracy of the Observatory, it will often be necessary to perform a peak-up before starting a spectroscopic integration in order to fine tune the telescope's pointing. Second, if the peak-up is done on a science target, it allows an observer to obtain a spectrum of a source even if its coordinates are not precisely known — as long as it is known to be bright enough for the peak-up to function properly.

There are three peak-up options available: **IRS Peak-Up**, **PCRS Peak-Up**, and **No Peak-Up**.

7.2.4.1 IRS Peak-Up

The IRS Peak-Up option uses software resident in the Combined Electronics to accurately find a target on one of the two IRS peak-up fields-of-view (blue or red) that share the SL detector array. The position determined for the peak-up target is then used to offset the science target to one of the IRS slits. An IRS peak-up can be performed on either the spectroscopic target itself, or a nearby ($\Delta R < 30'$) object whose offset from the science target is accurately known. If an offset peak-up target is desired, the observer can enter this directly into the AOT (see Figure 7.42 and detailed AOT Examples in section 7.2.7). See section 7.2.3.4.2 for a discussion on peak-up image quality.

The peak-up acquisition images are not supposed to provide science-quality data. However, some users may be interested in using the images to plan future observations. An approximate photometric calibration is available in the IRS Data Handbook and the SSC webpages.

7.2.4.1.1   IRS Peak-Up Filters

The blue IRS peak-up filter covers the 13–18.5 μm wavelength range, while the red IRS peak-up filter covers the 18.5–26 μm wavelength range. Users may select only one peak-up filter for a particular AOR. Figure 7.38 shows the filter transmission curves for the blue and red peak-up channels. Note that these curves represent only the filter transmission. They do not include the slightly varying responsivity of the detector with wavelength.

7.2.4.1.2   IRS Peak-Up Accuracy

Besides the peak-up filter, the observer must select one of three levels of accuracy for the peak-up, "**High,**" "**Moderate**," or "**Low**." These give positional uncertainties in placing and holding the science target on the slit of 0.4″, 1.0″, and 2.0″ (1σ radial).

Observers should note that the nominal initial pointing accuracy of Spitzer is ≤1.0″. The use of low-accuracy peak-up (1σ=2″) may degrade this performance and it is not recommended. In terms of time spent, the difference between High and Low accuracy peak-up is small; for a blue peak-up of a 5 mJy source, the AOR will be only 15 seconds larger than with Low accuracy peak-up.
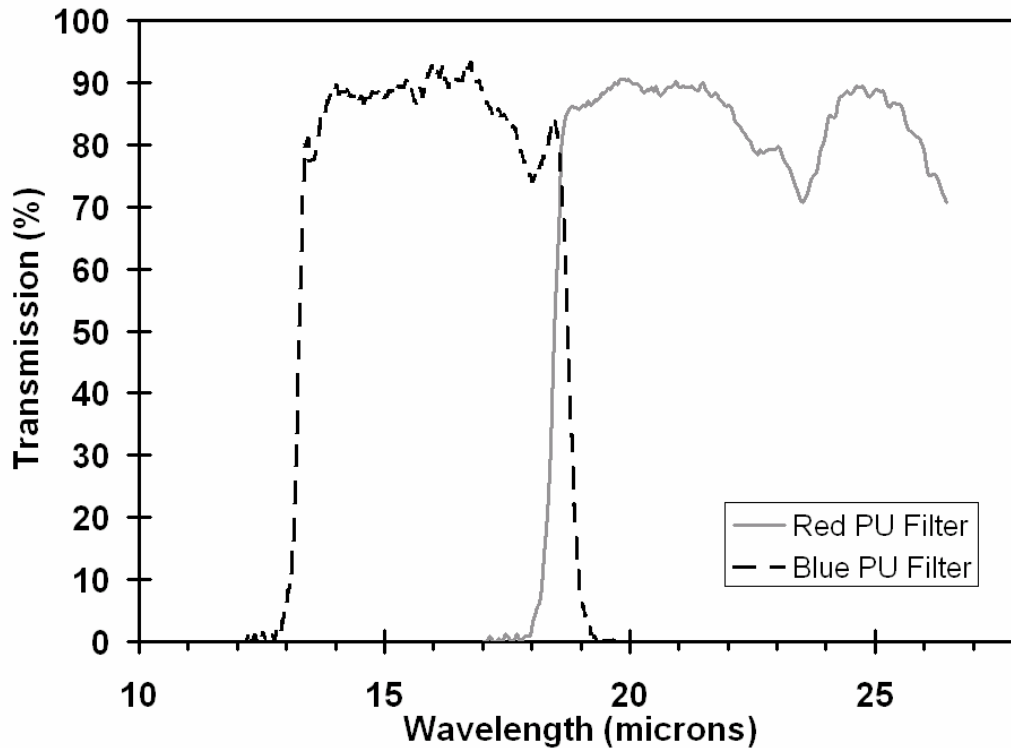
Figure 7.38: Transmission curves of the blue and red IRS peak-up filters.

Table 7.8 gives the photometric uncertainties corresponding to High, Moderate, and Low accuracy. These uncertainties arise because of the uncertainty in the fraction of source flux that enters the slit. Given the substantial overheads associated with peak-ups, observers should think carefully about the absolute flux uncertainties required for a given science program, and the loss in source flux that can be tolerated, in order to avoid specifying higher accuracy pointing than required. High accuracy pointing would be difficult to justify in a wide slit, for example, especially considering that the absolute flux uncertainties given in Table 7.8 in that case are comparable to or lower than the requirement on the science data pipeline processing. Conversely, using no peak-up with the narrow-slit SL or SH modules could result in the loss of substantial source flux. Use of the Moderate accuracy peak-up for observations that *require* a pointing accuracy of 1″ (in order to obtain the associated photometric uncertainty) is still advisable. This is especially true when multiple IRS modules are selected for use.

Table 7.8: Initial photometric uncertainties (95% confidence) for IRS Peak-Up options

	High	Moderate	Low
Positioning accuracy (arcsec), 1σ, radial	0.4	1.0	2.0
Photometric uncertainty (SL)	15–20%	37–42%	~100%
Photometric uncertainty (SH)	5–10%	15–20%	45–55%
Photometric uncertainty (LL & LH)	2–5%	5–10%	15–20%

7.2.4.1.3   IRS Peak-Up Target Brightness

If an IRS peak-up is specified, the peak-up target flux density (in the wavelength range of the selected peak-up filter) must be entered into the AOT. This flux density is used to select from a limited set of exposure times (see section 7.2.4.1.6). Entering an incorrect flux density would invoke non-optimal exposure times, and as a result the peak-up algorithm might fail to locate the peak-up source or waste unneeded exposure time on a bright source. **Peak-up targets must be brighter than 5 mJy (blue) or 15 mJy (red), and fainter than 150 mJy (blue) or 340 mJy (red).**

While the peak-up algorithm is capable of *acquiring* a point source as weak as 0.8 mJy in the blue or 1 mJy in the red, it does not meet the High accuracy requirements on such sources, having success rates of only 76% and 88% for red and blue, respectively (see section 7.2.4.1.7). On the other hand, the Medium accuracy requirement is met for the entire range of allowed source fluxes (for both blue and red peak-up modes), including the faintest sources.

> **It is important to note that the effective minimum brightness for IRS peak-up targets is a function of background brightness and will generally be *larger* than the absolute minimums discussed above — see Section 7.2.4.1.3.1 for additional details.**

In general, the performance of the peak-up algorithm is better when using blue filter observations than it is when using red filter observations. A peak-up using the blue filter allows one to acquire sources fainter than is possible with the red. Also, for a faint soruce, a blue peak-up entails much less integration time than a red one. The blue peak-up also has better centroiding precision for bright sources (however, under normal operating conditions for very bright peak-up targets, the red peak-up has a maximum flux limit higher than that of the blue peak-up). Spot will allow observers to enter peak-up targets with flux exceeding these values, but at the current time **we do not recommend peaking up on sources that will saturate the detector.**

7.2.4.1.3.1   Effect of Background on Peak-Up Target Selection

In some regions of the sky, the peak-up algorithm will not be able to acquire a point source with an acceptable rate of success because of the complicated background cirrus texture in the peak-up images. The point source flux must be greater than the confusion noise in the chosen ~1.0′ peak-up aperture for a

reasonable chance of acquisition. While most of the flux in the peak-up filter bandpasses comes from zodiacal emission, the zodiacal light is very smooth on arcminute scales, and we expect most of the texture to be due to cirrus and other Galactic backgrounds. To calculate quantitative lower limits on recommended peak-up point source fluxes, we use the cirrus sky power spectrum at 100 μm from Gautier et al. (1992, AJ, 103, 1313) to obtain the 100 μm confusion noise in the large-telescope (negligibly small PSF) limit for a 1.0′ aperture. We then place an upper bound on the 15 and 25 μm confusion noise in such an aperture by multiplying by an upper bound on the 12/100 and 25/100 flux density ratios obtained by Desert, Boulanger, & Puget (1990, A&A, 237, 215). Figure 7.39 shows the recommended minimum flux density of the peak-up target (in mJy) as a function of the cirrus brightness at 100 μm (in MJy/sr), for both the blue and red peak-up filters.
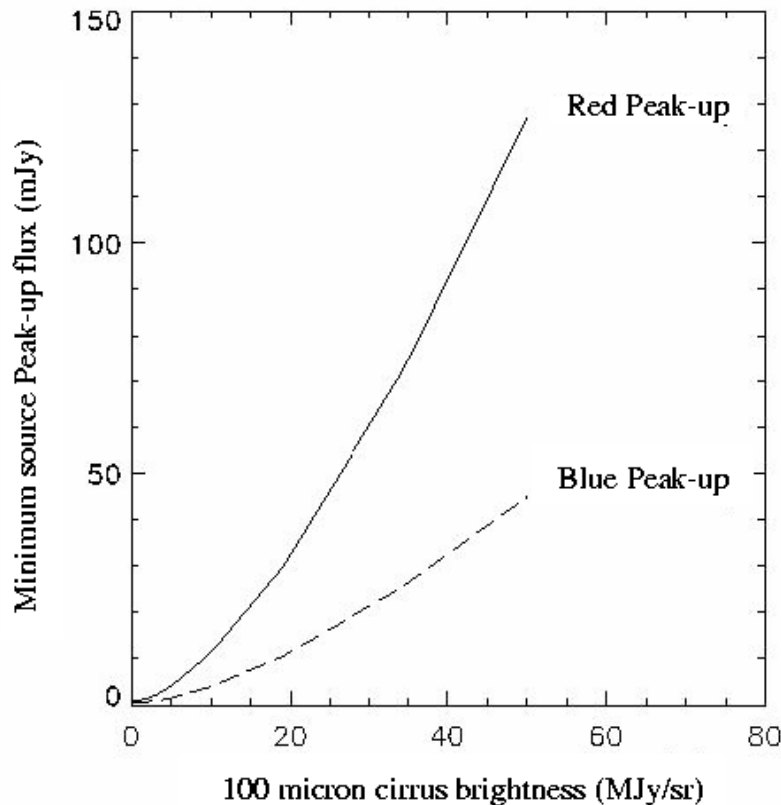
Figure 7.39: Recommended minimum flux density of the peak-up target f$_{blue}$(min) and f$_{red}$(min), in mJy, as a function of the background cirrus brightness at 100 μm (in MJy/sr).

7.2.4.1.4   *IRS Peak-Up Logic for Single and Cluster Targets*

The manner in which the selected High, Moderate, or Low accuracy pointing specified in a given AOR is achieved is essentially identical for both single targets and cluster targets. (A more complex scheme for the peak-up logic of cluster targets was described in earlier versions of the SOM.)

For a single target, no matter how many modules are requested in the AOR, a single peak-up is done at the beginning of the sequence. Multi-module observations are always scheduled in the order of narrowest-to-widest slit. Currently, the accuracy of positioning the object does NOT degrade for each Observatory move. Before the second day of IRS campaign 16 (9 December 2004), the accuracy of positioning the object would degrade by approximately 0.2″ added in quadrature for each Observatory move.

For cluster targets, a single peak-up is performed at the start of the observation, and the requested accuracy is maintained throughout the observation without intermediate peak-ups. If multiple modules are requested in the same AOR, then all cluster targets are observed with a single module first, before moving to another module. Peak-up is done at the beginning of the AOR.

7.2.4.1.5   Extended Sources as IRS Peak-Up Targets

When using the IRS Peak-Up, the observer is able to select either the **point source** (nominal) or **extended source** option. (This choice is not available for the PCRS Peak-Up.) The point/extended source selection sets important parameters in the peak-up software that are required to accurately centroid on an extended source. A point source is defined as having maximum dimension much less than the FWHM of the central Airy pattern of the telescope (~5″ with the red peak-up filter) – the majority of extragalactic sources and certainly all stellar sources fall into this category. An extended source is defined as having FWHM greater than the telescope Airy pattern, but not exceeding ~20″. The extended option allows for peak-up on cometary or comet-like sources with intensity distributions proportional to $1/R$ (where R is the radial distance from the source). **The recommended lower limit to the extended source peak intensity is 15 MJy/sr for the blue and 10 MJy/sr for the red, and the upper limits are 340 MJy/sr for both the blue and red, which is the saturation limit of the detector under normal operation.**

> **The IRS Peak-Up algorithm has been optimized for point sources. Therefore, we strongly recommend that, if at all possible, peak-up be performed on a point source, ideally on a relatively strong source in a region with relatively faint and unstructured background emission.**

For example, if the proposed scientific observation is of an extended target, then rather than peaking up on the extended target itself, the recommended procedure is to perform an offset peak-up on a nearby point source. The extended source peak-up option is to be used cautiously and only: (1) if there are no point sources with sufficient flux within $\Delta R < 30'$ of an extended target, or (2) if the position of an extended target is not well known, or (3) if the extended object is a moving target, such as a comet. It is also conceivable that peak-up be used to target the brightest sub-clump (FWHM < ~9″) of an extended source with complex morphology (using either the blue or red point source option), but this is risky and not recommended.

Peak-up centroiding accuracies on extended sources are poorer than those obtained with point sources. The performance of the peak-up algorithm was tested on simulated images composed of a single peak-up object on an otherwise flat background, with added CR hits, astronomical/detector noise components, and a constant zodiacal background. The "extended" sources were modeled with Gaussian intensity distributions, and cometary sources with $1/R$ intensity distributions. The peak-up accuracies depend, in large part, on the source size; centroiding accuracies progressively decrease with increasing source size. The maximal recommended FWHM of an extended peak-up target is ~20″. With such a target, one can expect centroiding accuracies of no worse than 1.5″ (with greater than 95% confidence levels) for peak intensities >15 MJy/sr (blue). The same numbers hold for cometary ($I \propto 1/R$) type sources.

7.2.4.1.6   *Operation of the IRS Peak-Up Algorithm*

The peak-up algorithm assumes that the source with the highest *peak flux* in the peak-up array field-of-view is the intended peak-up target. The field-of-view for peak-up acquisition has been reduced to 24×24 pixels, to avoid picking up neighboring bright sources. The algorithm for determining the peak flux, and the area over which the centroiding is performed, differs between the point- and extended-source cases:

- Point source option: The peak flux of a given source is the sum within a 3×3 pixel (5.4″×5.4″) box centered on the peak pixel of the source. Centroiding is performed within a 5×5 pixel (9″×9″) box centered on the same peak pixel.
- Extended source option: The peak flux of the source is the sum within a 7×7 pixel (12.6″×12.6″) box centered on the peak pixel of the source. Centroiding is performed within an 11×11 pixel (19.8″×19.8″) box centered on the same peak pixel.

The peak-up algorithm utilizes a set of three DCS images, acquired in sequence, containing the peak-up target specified by the observer. To identify and accurately centroid on the peak-up target, the algorithm does the following:

1. Produces three difference images, (DCS1 - DCS2), (DCS2 - DCS3), and (DCS1 - DCS3), which presumably contain only noise and transient signatures (such as cosmic rays). The algorithm then makes a rough identification of cosmic ray (CR) hits under the assumption that the strong CR hits produce the upper and lower 10% of outlier values in the histograms of the difference images.
2. Using the CR hits identified in step 1, isolates the noise component in the difference images, replaces bad pixel values, and estimates the noise variance (σ).
3. Removes any planar background component via a least-squares approach and then subtracts the mean from the three original DCS images, producing a set of zero-mean images.

4. Identifies as CR hits (or transients) those pixels with values outside of $\pm 3.5\sigma$ in the difference images, then replaces those pixel positions in the zero-mean images with zero values.
5. Matches filters by convolving the appropriate (blue or red) point spread function with the three zero-mean images, and then co-adds them into a single image.
6. In the co-added image, finds isolated pixel clusters below a minimal pixel size threshold (containing less than 6 pixels for red or 15 for blue) and then assigns zero values to those pixels. These are assumed to be CR hits, noise spikes, or glitches that have somehow escaped earlier removal.
7. Centroids on the source(s) in the final co-added image and reports the centroid of the source with the greatest peak flux to the telescope computer. A final co-added image is returned to the ground.

The acquisition of the three DCS images and the subsequent processing by the peak-up algorithm is actually done in two stages whenever an IRS peak-up on a source is requested. The first stage (referred to as the "acquisition" peak-up) acquires the source and then requests a telescope move to place the peak-up target on a portion of the peak-up array least affected by cosmetic defects (known as the "sweet spot"). The second stage (referred to as the "sweet spot" peak-up) finds the centroid of the source in the sweet spot region before offsetting the target to the slit.

Consequently, each IRS peak-up produces six, unprocessed differenced DCS images, as well as two dark-subtracted, cosmic-ray cleaned, flat-fielded, and background-subtracted "average" frames. There are three DCS images for the acquisition peak-up and three DCS images for the sweet spot peak-up. Each of these, in turn, produces a single, processed, average frame. These average frames are used by the algorithm to find the peak-up target. Thus, the observer selecting an IRS peak-up will receive, in addition to the desired spectral data, the set of eight peak-up images produced by the peak-up algorithm. These eight images are not passed through the IRS science data processing pipeline.

7.2.4.1.7   IRS Peak-Up Performance Characteristics

Table 7.9 details the performance of the peak-up algorithm when applied to a point source on a flat background, with no other sources within the field. The test trials were performed using 1024 simulated DCS image three-sets, each three-set containing a single point source Airy pattern randomly located in the field. Also included in the trial images were simulated cosmic ray hits, the standard noise components (readout noise, dark current, and photocurrent shot noise) and a constant offset level representing the zodiacal light component. In all simulated images we applied a background level equivalent to the zodiacal component at $30°$ ecliptic latitude, and CR hits at a rate of 8 cm^{-2} s^{-1}. These are the approximate average observing latitude and 95% confidence level CR hit rate expected for Spitzer. In-orbit measurements show a CR hit rate of ≈ 2 cm^{-2} s^{-1}. Table 7.9 lists the success rates with respect to the High and Medium accuracy pointing

requirements, as a function of the point source flux and peak-up filter used (blue or red), for three ecliptic latitudes. The "success rates" are defined as follows:

- **High accuracy:** The percentage of trials in which the determined source centroid is within $\Delta R = 0.42''$ of the actual source position. Peak-up fulfills this requirement, for a source with flux F, if it meets the above constraint with a success rate of 95% or greater.
- **Medium accuracy:** The percentage of trials in which the determined source centroid is within $\Delta R = 0.84''$ of the actual source position. Peak-up fulfills this requirement, for a source with flux density F, if it meets the above constraint with a success rate of 95% or greater.

Table 7.9: IRS Peak-Up success rates from simulations

	Source Flux Density (mJy)			Success Rates	
	Ecliptic plane	Ecliptic latitude=30 °	Ecliptic pole	High	Medium
Blue	$1.1 < F < 1.4$	$0.8 < F < 1.0$	…	88 – 95 %	> 97 %
	$1.4 < F < 150$	$1.0 < F < 150$	$0.8 < F < 150$	95 – 99 %	> 99 %
Red	$1.4 < F < 7.0$	$1.0 < F < 5.0$	$1.0 < F < 3.6$	76 – 96 %	> 98 %
	$7.0 < F < 340$	$5.0 < F < 340$	$3.6 < F < 340$	98 – 99 %	> 98 %

In-orbit testing shows results consistent with those discussed here and listed in Table 7.9. The success rates are based on the integration times used in flight, which, in turn, depend on the source flux: (i) $T_{acq} = T_{sweet} = 8$ seconds per single DCS frame for both the blue and red filters for sources with F > 5 mJy, and (ii) T_{acq} and T_{sweet} integration times ranging between 8 and 64 seconds per single DCS frame for sources with $F < 5$ mJy.

If the peak-up algorithm fails (also see Section 7.2.4.1.10), then the AOR will still be executed and the slit(s) will be positioned on the science target(s) at the nominal pointing accuracy of Spitzer (1.0″).

The point source peak-up results presented thus far are based on idealized scenarios in which there are no complex background components such as faint extended clumps or cirrus emission, and are consistent with in-orbit results. The mid-infrared background is too diverse, and its structure too poorly known, for comprehensive characterizations of its influence on peak-up success statistics. The results in Table 7.9 should thus be considered optimistic estimates, perhaps best reflecting extragalactic observing conditions in which the science target is point-like, is likely the only source in the field, and the field contains little structured background emission. As noted in section 7.2.4.1.3.1, in some regions of the sky, the peak-up algorithm will not be able to acquire a point source with an acceptable rate of success because of the complicated cirrus texture in the peak-up images.

7.2.4.1.8   IRS Peak-Up Target Selection for Stationary Science Targets

Spot features an algorithm that automatically selects IRS peak-up stars from the Two-Micron All-Sky Survey (2MASS). These stars have well-determined positions, to better than 0.7″. Their near-infrared 2MASS photometry can yield predictions for the mid-IR flux densities in the IRS blue and red peak-up filters using stellar photospheric models.

The Spot algorithm needs, as input, the target equatorial coordinates (J2000.0). The output is a table of stars from 2MASS that are suitable for peak-up in either the red or blue filter. The user selects one of the stars from the table as the IRS peak-up star. The table lists, for each star, R.A. and DEC (J2000.0), coordinate epoch, derived stellar spectral type, and flux densities and uncertainties (mJy) in the blue and red peak-up filters. The table also indicates if the star is readily "GOOD for Blue" or "GOOD for Red" peak-up, or merely "Acceptable." The latter might have infrared excesses from dust shells, which could saturate the peak-up arrays, although such cases cannot be firmly determined a-priori.

7.2.4.1.8.1   Operation of the 2MASS Peak-Up Search Routine

- The search routine is executed in Spot via the "2MASS selection" button found under "Peak-Up Settings" in the Staring or Spectral Mapping AOT entry windows.
- Before clicking on "2MASS selection," make sure you have chosen the desired peak-up band. The correct flux will then be entered in the Peak-Up Target field after a star is selected from the pop-up table.
- The "Surface Brightness" pop-up window will appear. The local background surface brightness at 100 μm is needed by the 2MASS peak-up star search routine. This parameter is used to determine the minimum peak-up star flux density (always ≥ 10 mJy) — see section 7.2.4.1.3.1. The surface brightness can either be entered directly or estimated by Spot. If the surface brightness is > 50 MJy/sr, then an error message will be issued by Spot, as structure in extended background emission brighter than this is expected to interfere with the peak-up algorithm. **Under no circumstances should the "OK" button be pressed with the surface brightness value still set to the default of 0.00 MJy/sr.** Doing so results in the minimum peak-up star flux density being set to 10 mJy, which may return many peak-up star candidates that are not appropriate for the true (non-zero) background surface brightness. Selecting one of these stars could result in a failed peak-up.
- If no adequate 2MASS star is found, a pop-up window will read, "No 2MASS object found as IRS peak-up candidate."
- If more than one 2MASS peak-up star is found and displayed in the pop-up table, then the observer should select (in priority order) "GOOD for Blue" stars, then "GOOD for Red" stars, then "Acceptable" stars (only if no "GOOD" stars are available). Among "Acceptable" stars, preference should be given to earlier spectral types, and/or fainter flux densities at 16

μm.

- In order for the peak-up to succeed, there can be no bright objects within 1′ of the peak-up star. Spot filters out candidate peak-up stars where this is the case. However, observers should still visually check their peak-up stars to confirm this, in particular, that a much brighter object than the peak-up star is not within this radius.

Spot derives the table of suitable peak-up stars as follows. First, stars in the 2MASS All-Sky Data Release Point Source Catalog that are free of photometric defects or contaminating artifacts (as indicated by the 2MASS data quality flags), are extracted within a circle of 30′ radius centered on the science target.

Second, using the 2MASS $(J–K_s)$ and $(H–K_s)$ colors, the stellar spectral type and luminosity class of the 2MASS sources are estimated. The spectral type is needed because the flux densities of stars in the IRS peak-up filters, relative to near-infrared flux densities, vary by 60% from B to M0 stars (Figure 7.40). The luminosity class is needed to identify late-type giant and supergiant stars. These sources are likely to have circumstellar dust shells that would cause the predicted mid-infrared flux densities (see next step) to underestimate their actual fluxes. Some of these stars are nonetheless output as "Acceptable" for peak-up. The algorithm to determine type and class is based on theoretical stellar models of $(H–K_s)$ vs. $(J–K_s)$ colors for dwarfs, giants, and supergiants by Koornneef (1983, A&A, 128, 84). Spot currently assumes that 2MASS colors are equivalent to Koornneef colors; however, this may be updated in the future to use only fiducial 2MASS colors for estimating the corresponding spectral types. Given the 2MASS star's position in the $(J–K_s)$, $(H–K_s)$ color-color diagram, the algorithm determines which of the loci (for dwarfs, giants, or supergiants) intercept the rectangular region centered on the star's $(J–K_s)$ and $(H–K_s)$ colors and bounded by $\pm2.5\,\sigma_{(J-K_s)}$ and $\pm2.5\,\sigma_{(H-K_s)}$ if $(J–K_s) < +0.83$ or $(J–K_s) > +0.93$. Otherwise, it is bounded by $\pm\sigma_{(J-K_s)}$ and $\pm\sigma_{(H-K_s)}$. The loci that intercept this rectangular region indicate the possible luminosity classes of the 2MASS star. The algorithm then determines, separately for each possible luminosity class, the spectral type from the $(J–K_s)$ color, using the theoretical relations by Koornneef (1983). The final spectral type is the mean of those from the possible luminosity classes. It is expressed as a range, given the uncertainties $\sigma_{(J-K_s)}$ and $\sigma_{(H-K_s)}$, and the various possible luminosity classes indicated for the star.

Third, photospheric flux densities in the peak-up filters are predicted from the 2MASS K_s magnitude and the estimated spectral type and luminosity class. The 2MASS K_s magnitude has absolute calibration $F_v(2.17\ \mu m) \approx 667 \times 10^{-0.4K_s}$ (Jy). Then, the stellar spectral models of Lejeune et al. (1997, A&AS, 125, 229) are utilized; these models are based on Kurucz models but improved for late-type stars. These models yield flux density ratios at 16 or 22 μm (the approximate values for the IRS blue or red peak-up filter effective wavelengths, respectively) compared to 2.17 μm. Figure 7.40 shows such ratios for dwarfs, giants, and supergiants of various spectral types. Given the type and class from step 2 above,

and the flux density F_ν (2.17 μm), Figure 7.40 is used to derive F_ν (16 μm) and F_ν (22 μm). The uncertainties in these flux densities arise from those in the spectral type and the various possible luminosity classes from step 2 above.
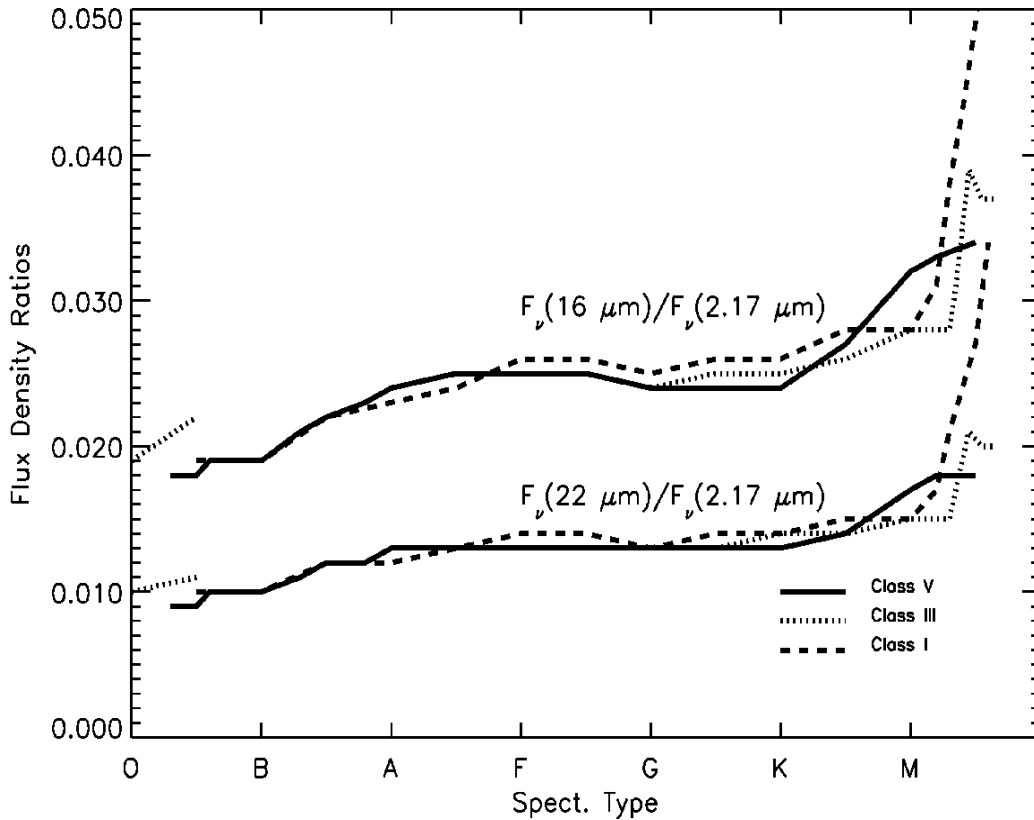
Figure 7.40: Flux density ratios of stellar photospheres for different luminosity classes, between the IRS peak-up filters (approximate central wavelengths of 16 μm for "blue" and 22 μm for "red") and the 2MASS K_s (2.17 μm) filter. Based on Kurucz models by Lejeune et al. (1997, A&AS, 125, 229).

Finally, if f_{blue}(min) < F_ν (16 μm) < 150 mJy, then the star is a candidate for a blue filter peak-up. If f_{red}(min) < F_ν (22 μm) < 340 mJy, then the star is also a candidate for a red filter peak-up. The lower (faint) flux limits f_{blue}(min) and f_{red}(min) are determined as a function of the 100- μm background surface brightness (see Figure 7.39). The Spot algorithm outputs the star as one of the following: (a) "GOOD for Blue" if F_ν (16 μm) satisfies the corresponding relation above and either (J–K_s) < +0.83 or the luminosity class is unambiguously "V" (dwarf); (b) "GOOD for Red" if the conditions in (a) are not met, but F_ν (22 μm) satisfies both the corresponding relation above and the remaining conditions from (a); or (c) "Acceptable" if (J–K_s) > +0.83 and the luminosity class is not "V" (in this case, the star is a giant/supergiant with a possible dust shell that may contaminate its long wavelength fluxes). Figure 7.40 shows that the flux ratios vary rapidly with spectral type for stars later than type M0. Therefore, the

predicted flux densities of late-type giants/supergiants must be treated with caution.

7.2.4.1.9   IRS Peak-Up for Moving Science Targets

The peak-up capabilities for moving targets are peak-up on the science target, peak-up on an offset target that is co-moving with the science target, and peak-up on an inertial target and the offset to the moving target. As with peak-up on a stationary target, the peak-up algorithm simply centroids on the brightest object in the peak-up array field-of-view; that is, the peak-up algorithm does not "know" that the selected peak-up target is a non-stationary object, nor does it discriminate between moving and stationary objects detected in the same peak-up field-of-view. Results from IOC and SV indicate that IRS peak-up was successful for all moving targets attempted, including objects with rates of motion between $<1.0''$/hr and just over $200''$/hr, and 15 μm fluxes between 40 mJy and 1 Jy. Both moving point sources and moving extended sources were acquired successfully.

To take spectra of a single moving target, and also peak-up on that target, enter the target as a "moving single." Then, in the IRS AOT, set the peak-up offsets to zero in both R.A. and Dec. Note that High accuracy is unlikely to be achievable for tracked observations of moving targets, given the errors anticipated during tracking. The observer can still specify a High accuracy peak-up for a tracked moving target, which will result in the most accurate initial peak-up. However, it will not deliver the same accuracy of pointing or stability anticipated for a fixed target. The actual performance for moving targets will continue to be refined as in-orbit data are accumulated and analyzed; updates will be posted on the SSC website. If spectrophotometry of a moving target is required, this can probably be most robustly achieved with a short spectral map, rather than specifying a High accuracy peak-up in staring mode. This spectral map can be requested with or without a peak-up prior to the scan, depending on the anticipated accuracy of the object's ephemeris.

In the case where you want to observe a single moving target but peak-up on a co-moving offset source, enter the science target as above, but enter the co-moving peak-up source position in the peak-up target offset entry box. If, instead, you want to peak-up on a moving target (for which you have a NAIF ID) and then observe one or more co-moving positions, enter this target as a "moving cluster" with the target for which you have the NAIF ID as the "primary target" and the offset(s) you wish to observe as the cluster position(s). Here, the peak-up will be performed on the primary target, and spectra will be obtained at the primary and offset positions. If you wish to peak-up on the primary target but not obtain a spectrum of it, simply check the "observe offsets only" box in the moving cluster target entry window.

7.2.4.1.10 Assessing IRS Peak-Up Success

The following information can be used by observers as a guide for assessing the success of an IRS Peak-Up after receiving IRS data.

7.2.4.1.10.1 What to Look for From a Successful IRS Peak-Up

At each of the two peak-up centroiding positions – acqusition (ACQ) and sweet spot (SS) – three images (DCEs) are taken and processed onboard Spitzer to produce the image frames that are used by the peak-up algorithm. This fourth, processed frame is always DCE number 0003, and the exposure number is 0000 for ACQ and 0001 for SS. Hence, to examine the peak-up data, the first step is to look at the files with exposure.dce numbers of 0000.0003 and 0001.0003.

The FITS headers for the peak-up images contain information about the operation of the telescope and the peak-up algorithm (see below). The red and blue peak-up arrays share a common WCS that is provided in the FITS header. The WCS is described in the system CTYPE1 = 'RA---TAN-SIP', CTYPE2 = 'DEC--TAN-SIP', not 'RA--TAN' and 'DEC--TAN', respectively. The data in each FITS file show the combination of the three individual exposures, including processing for cosmic ray rejection, flat fielding, and background subtraction. The data are in units of DN. To convert to electrons per second, multiply by 4.6 to get to electrons, and then divide by the exposure time given in the RAMPTIME header keyword.

In the processed images, any source that is bright enough for centroiding will be clearly visible. Assuming that the predicted flux of the peak-up target was correct, the peak-up target will be located near the center of the image. In the ACQ frame, it will be at or near pixel (107, 30) for blue or (105, 92) for red; in the SS frame, it will be at pixel (108, 28) or (106, 94) for blue or red, respectively. The onboard software considers the lower left corner of the array to be (0, 0) with the lower left corner of that pixel (-0.5,-0.5); that is, pixel centers are labeled with integer coordinates.

It is possible that other sources will be visible in the field, as well. It is the responsibility of the observer to have vetted the peak-up field to ensure that the brightest source is the intended peak-up target (see section 7.2.4.6). The brightest source in each of the ACQ and SS frames will be selected as the peak-up target by the peak-up algorithm.

Note that with the current onboard software processing, column-by-column variation may be present in the image. This is a known effect (informally referred to as "jail-barring"). The peak-up algorithm has been demonstrated to work as advertised in the presence of this effect. Future revisions of the onboard software will mitigate against this effect.

7.2.4.1.10.2 Spotting a Failed IRS Peak-Up

There are two ways for an IRS peak-up to "fail."

1. If the algorithm does not find a valid centroid, then it will report a failure. In the peak-up images, particularly the processed DCE (see above), there will be no visible source. These failures can also be identified from the FITS header by looking at the value of the "peak-up centroid quality

code," which is 0 for a failure and 1 for a success.

2. The peak-up algorithm can also result in a false positive (also see section 7.2.4.1.10.3). In this case, the peak-up software centroids on something other than the intended target and reports a success. The peak-up images can be examined to see if the intended peak-up target is at the centroid reported in the header. The value of the centroid is given in the AXCNTRD1 and AYCNTRD1 keywords; note that these values are in centipixels, so need to be divided by 100.0 to compare with pixel positions in the image.

In addition, the value of the PTGDIFF keyword in the header gives the difference in arcseconds between the requested and reconstructed pointing. This value is almost always less than 1 arcsecond, which indicates that the intended coordinates were placed on either the acquisition or sweet spot.

7.2.4.1.10.3 IRS Peak-Up False Positives

The integrated number counts for sources with flux density of 1 mJy at 24 μm at high latitudes (see Marleau et al. 2004, ApJS, 154, 66) give 335 sources (both stars and galaxies) per square degree in a high galactic latitude field. That is one source per 10 peak-up fields. At a 24-μm flux density of 5 mJy, the number counts are only 25 per square degree or one source per ~150 peak-up fields. So, for a peak-up source with an estimated flux of 1 mJy, it is not that unusual to have a second source in the peak-up array field-of-view that is as bright or brighter. The source density can be much higher in many regions of the sky. If the selected peak-up target is near the faint flux density limit for the IRS peak-up arrays (5–10 mJy), then the peak-up data should be carefully examined to ensure that the peak-up was performed as intended and, therefore, that the IRS spectrum was obtained for the correct target.

7.2.4.2 PCRS Peak-Up

The second peak-up target acquisition option is PCRS Peak-Up. This choice selects the Pointing Calibration and Reference Sensor (PCRS) as the peak-up instrument. If PCRS Peak-Up is selected the user will not receive any peak-up images. The PCRS operates in the visual part of the spectrum (from 5050–5950 Å), and its main function is to calibrate and remove the drift between the star trackers and the telescope. The field-of-view for the PCRS is $40''\times40''$, comprised of a 4×4 array of $10''$ pixels. Consequently, the PCRS functions more like a "quad cell" centroiding device than a standard CCD that adequately samples the PSF. The PCRS can measure the centroid of stars in the 7.0 mag < V < 10 mag range to an accuracy of $0.14''$ (1σ radial). As in the case of the IRS Peak-Up for a stationary science target, the observer has the choice of peaking up on either the science target itself or an offset star.

Given the large wavelength difference between the PCRS and the IRS detectors, it is **strongly recommended that the PCRS only be used to peak-up on an offset stellar source with a known visual magnitude**.

The primary reference catalog for the PCRS is the PCRS Guide Star Catalog, which contains a carefully selected subset of stars from the Tycho catalog. In the future, it will be possible to select PCRS Guide Star Catalog stars from within Spot; meanwhile, the PCRS Guide Star Catalog and an IDL-based search routine are available on the SSC website.

Observers who choose to specify a PCRS peak-up star that is not listed in the PCRS Guide Star Catalog should consider the following conditions when selecting the star:

- It must be in the magnitude range 7.0 mag < V < 10 mag.
- There should be no neighbors within 1 arcminute that are brighter than 5.5 magnitudes fainter than the peak-up star.
- It should not be a variable star.

Any star selected for use in a PCRS peak-up, whether from the PCRS Guide Star Catalog or another source, should be visually inspected (e.g., using the Digitized Sky Survey) for the presence of nearby neighbors and/or overlapping, bright, extended emission. PCRS peak-ups are currently treated as High accuracy peak-ups. In the future, it will be possible to select a Moderate accuracy PCRS peak-up mode that will invoke peak-up on a visual target without incurring unnecessary peak-up and overhead burdens. Selection of peak-up stars and use of the PCRS Guide Star Catalog is discussed in more detail in the Spitzer Observation Planning Cookbook (which is available on the SSC website).

7.2.4.3 No Peak-Up

The third peak-up target acquisition option is No Peak-Up. In this case, the telescope will slew to the observer-specified science target position without refining the pointing accuracy, wait for the nominal settling time, and begin the spectroscopic observations. Positional uncertainty in this case is the nominal pointing accuracy of Spitzer, 1.0″.

7.2.4.4 Summary of Recommendations for Peak-Up on Point Sources

- Use the blue peak-up filter when possible (unless the peak-up target is very much weaker in the blue than in the red). A blue peak-up typically gives better centroiding accuracies than a red one.
- In general, supplying a peak-up target flux in Spot that is accurate to within a factor of two of the actual flux will improve the noise statistics and, hence, the probability of success for achieving the desired centroid accuracy.
- For sources brighter than 5 mJy, the peak-up duration time is expected to be the same regardless of the accuracy requested. In such cases it seems prudent to request a High accuracy peak-up (but be aware of possible additional overheads that might result from asking for many High

accuracy peak-ups in a single AOR).

- If at all possible, use a point source with flux in the midrange of f_{blue}(min) to 150 mJy for the blue, or f_{red}(min) to 340 mJy for the red. High accuracy peak-ups on weak sources (< 1 mJy for the blue and < 5 mJy for the red) have poorer success rates (below the requirement level of 95%) and require longer integration times. Moreover, the faint source peak-up success rates suffer the most rapid degradation when structured background emission is present.
- Choose a peak-up target that is a point source in a region with as little structured background emission as possible. To ensure a successful peak-up, avoid regions where the intensity of any structured background component (within the size scale of the ~1′ peak-up field) is greater than 10% that of the peak-up target. Refer to Figure 7.39 for a guide to the minimum recommended peak-up point source flux as a function of the 100-μm background cirrus brightness.
- Use Spot to select IRS peak-up stars from 2MASS.

7.2.4.5 Summary of Recommendations for Peak-Up on Extended Sources

- Limit extended peak-up targets to a size of FWHM < 20″.
- If peaking-up on a comet, use the "extended" source peak-up option.
- Peak-up on simple sources. Avoid sources with prominent and/or multiple sub-clumps.
- Choose a field as clear as possible of structured background emission. In the case of an extended source peak-up, the algorithm uses a 20″×20″ centroid box. All emission within that box will influence the centroiding results.
- For extended/cometary sources with peak intensities of 15 MJy/sr to 340 MJy/sr (blue) or 10 MJy/sr to 340 MJy/sr (red), peak-up should provide centroiding accuracies of < 1.5″, at a 95% confidence level.

7.2.4.6 Summary of Peak-Up Configuration Criteria

During the early mission, the SSC noted several instances in which the IRS peak-ups configured by observers did not yield the desired results (i.e., the peak-up algorithm did not center on the intended peak-up target). Peak-up configurations (both IRS and PCRS) in AORs should meet the recommended criteria listed below.

7.2.4.6.1 All Targets

- The peak-up target must be within 30′ of the science target.
- Point sources are recommended for use as peak-up targets. If using an extended source for a peak-up target, it must have a centrally peaked flux distribution and FWHM smaller than 20″.
- Blind pointing of Spitzer is much better than anticipated before launch. Intrinsic pointing accuracies of 1–2″ or better are routinely achieved.

Therefore, all observers should carefully assess their photometric accuracy requirements and peak-up needs. For example, if the target coordinates are accurately known, and a wide slit is being used and/or the highest level of photometric flux calibration accuracy is not required, then use of a peak-up may not be needed at all.

- The peak-up target must be the brightest object in the field-of-view of the selected peak-up array (~60″×80″ for IRS Peak-Up, ~40″×40″ for PCRS Peak-Up), in the wavelength band of the selected peak-up array (13.3–18.7 μm for IRS Blue, 18.5–26 μm for IRS Red, 5050–5950 Å for PCRS Peak-Up). Observers must manually inspect the fields around ALL peak-up targets (e.g., using Spot) to ensure that this condition is met. NOTE: since the optical, near-IR, and mid-IR colors of sources can be very different, ensuring that the peak-up field is "clean" at 2 μm (i.e., 2MASS K-band) is a necessary, but not sufficient, criterion for a successful peak-up. Whenever possible, inspect an image of the peak-up field obtained in a wavelength band as close as possible to that of the selected peak-up mode.

- Ensure that the best estimate of the peak-up target's IR flux density in the selected wavelength band (IRS Peak-Up) or V magnitude (PCRS Peak-Up) has been entered into the AOR. In addition, the brightness of the (point source) peak-up target should fall within the recommended brightness range for the selected peak-up mode: 1–150 mJy (IRS Blue), 5–340 mJy (IRS Red), or 7 < V < 10 (PCRS). In general, it is recommended that IRS peak-up targets have a minimum flux density of 10 mJy in the selected wavelength band (Blue or Red). The minimum allowed target brightness for a successful IRS Peak-Up depends on the local background surface brightness (see the Spitzer Observation Planning Cookbook for more information).

7.2.4.6.2 Stationary Targets

- Use the best known J2000 coordinates for the peak-up target (including cumulative corrections for proper motion since the epoch of coordinate measurement), or the best known offsets from the science target.
- Ensure that the best known proper motions for the peak-up target (and epoch of coordinate measurement) have been entered into the AOR. This is especially important for peak-ups on offset stars, but could also affect peak-ups on science targets with large proper motions.

7.2.4.6.3 Moving Targets

- Peak-ups must be configured on the moving science target (or a co-moving target), so the position of the target must be known to better than the dimensions of the selected peak-up array. Spitzer-centric positional error ellipses for the current ephemerides for objects likely to have large positional uncertainties (e.g., near-Earth objects, Kuiper belt objects, and comets) are available through JPL's Horizons database — see http://ssc.spitzer.caltech.edu/documents/horizons.pdf.

- Objects with large positional errors at the time of proposal submission can still be requested, but should be identified as such in the proposal. Note, however, that objects with large positional errors may not be scheduled if the ephemeris is not subsequently improved prior to the Spitzer observations. The SSC will endeavor to use the best ephemeris available from the Horizons database at the time of scheduling, typically 5 weeks before execution of the observation.
- Some targets have positional uncertainties that are initially relatively small, but are likely to *increase* in the 5-week interval between generation of the spacecraft schedule and execution of the observation. This type of behavior may be exhibited by objects that experience non-gravitational forces, such as comets nearing perihelion. In such cases, the observer may request a late ephemeris update.
- Typically, a moving target observer will not know the scheduled date of observation for his/her target at the time that the target is proposed. Consequently, he/she cannot check the corresponding peak-up array field-of-view for bright neighbors that could confuse the peak-up algorithm (see above). However, the SSC strongly encourages moving target observers to use the Spot visualization tool to check for bright objects along the orbital track corresponding to the two Spitzer visibility windows for the target during the proposal Cycle. This can be done by choosing a date in the middle of each visibility window, and overlaying the moving target orbital track on a background image closest in wavelength to the peak-up array selected for the Spitzer observations. Timing constraint(s) may have to be applied to the AOR to avoid scheduling the observation at a time when the peak-up could fail.

7.2.5 Additional Recommendations for Observation Planning

7.2.5.1 Spectrophotometric Errors in Short-Low Data Resulting from Saturated IRS Peak-Up Arrays

When the IRS arrays are non-destructively read in SUR mode, they suffer from "droop," which is the addition of a voltage to each pixel proportional to the charge summed over the entire array; see section 7.1.7. Droop is a detector artifact, not a true signal, and is present in all IRS arrays. An estimate of the total charge falling on the array must be made before a proper droop correction can be calculated, and this is done in the IRS science data processing pipeline (see Figure 7.55). However, the IRS peak-up apertures share the SL array with the 1st and 2nd orders and 1st order bonus spectra. The entire array is exposed during a SL exposure, and since the blue and red peak-up filters are broad, the peak-up fields can reach the A/D limit quite quickly, especially in regions of high surface brightness (for example, at low zodiacal or Galactic latitudes, or high brightness nebulae). However, *even in dark regions of the sky*, the peak-up arrays may saturate during the longest ramp duration available (240 s). When A/D saturation occurs in the peak-up arrays, it is much harder to predict the appropriate droop

correction for the rest of the array from that moment onwards, because many thousands of peak-up pixels are still accumulating charge, but are not being recorded. If the peak-up arrays saturate at any time during integration, the result will be an increased uncertainty in the shape of the spectral order ramps. The situation is worst when peak-up saturation occurs very early in the ramp (for example, if the peak-up arrays saturate by the second sample of the observations), but saturation *at any time* during the integration will lead to photometric uncertainty in the spectra.

Although, in principle, we can predict the charge accumulating on the peak-up arrays after A/D saturation, these predictions depend on modeling the behavior of the detectors above the A/D limit. Recent attempts to parameterize this shape are producing acceptable results (IRS S15 pipeline and onwards), but **we recommend that, whenever possible, observers should try to design their observing program in such a way as to avoid strong saturation of the peak-up array**. We note that, in some cases (for example, 240-s integrations on sources in otherwise low surface brightness backgrounds), it may be impossible to avoid peak-up saturation at some point in the accumulating exposure. In those cases, the observer must rely on the pipeline to correct for the significant droop effect. We recommend that, for observations requiring high photometric precision, the observer should check the shapes of ramps in their data in the period before and after peak-up saturation, to ensure that the slopes in these regions are consistent. **In most cases, however, it should be possible to design a set of observations that do not saturate the peak-up arrays at all. This is the safest mode of operation for the SL module.** It should be noted that point source saturation in the peak-up arrays, although subject to the same considerations as above, usually have less of an effect on the droop corrections, because this saturation affects a much smaller number of pixels. We consider below two cases (uniform and point source illumination limits for the peak-up apertures) to derive rules of thumb for avoiding problems resulting from this effect.

Extended source limits to avoid peak-up saturation: The last sample in the ramp for the peak-up arrays (at 16 and 22 μm) will saturate at 450 MJy/sr in 6 seconds. However, the current droop correction is appropriate for slightly larger limits: 500 MJy/sr for 6 sec, 200 MJy/sr for 14 sec, 50 MJy/sr for 60 sec, and 25 MJy/sr for 240 sec ramps. We note that even with the improved droop prediction software, observers should never exceed a surface brightness of 114 MJy/sr for a 240-sec integration, which is a hard limit for peak-up saturation before the second read of the array.

We recommend that you use Spot to check the expected background levels in the vicinity of your targets during the allowable observing windows, and ensure that you do not exceed the limits cited above for any of your SL observations.

Point source limit: The limits on point source saturation are less stringent than those on extended source saturation, because of the smaller number of affected pixels. To stay with a 4% radiometric accuracy, observers should avoid saturating

more than 180 pixels by the second sample of a single 240 s exposure. Based on our current best knowledge of the Spitzer PSF, this is produced by a point source of approximately 41 Jy, for either blue or red peak-up.

As mitigation, observers may reduce the ramp time (while increasing the number of cycles), or constrain the observation date as to avoid the bright region or source, if possible.

7.2.5.2 Dedicated Observations for Background Subtraction

The mid-infrared background, predominantly from zodiacal light at IRS wavelengths, can contribute significantly to observations of faint targets with the IRS. The arrays are sufficiently sensitive to detect levels of 20 MJy/sr or more at any wavelength. The determination of when to obtain separate (i.e., off-source) background observations, in the context of the information presented in this section, is left up to the observer. If the background is comparable in brightness to the science target, or if detection of features at the level of the background is attempted, then a proper background subtraction is critical. Therefore, observers are recommended to examine their data and use the background estimation tool in Spot to gauge the expected background levels in order to weigh these against the science goals.

We recommend careful background subtractions for all low-resolution (SL and LL) observations. This is especially true if the science goals critically depend upon measuring an accurate spectral continuum shape or absolute flux level for the target, or if the target is extended relative to the sizes of the low-resolution slits. Background subtraction should be performed on SL and LL data before making any absolute photometric scaling. In Staring mode, all science targets are placed at two nod positions along the IRS slits, so this background subtraction can be accomplished by differencing the nod positions. Alternatively, the accompanying subslit (e.g., SL2 for SL1 spectra) can also be used for background subtraction. It is important to realize that the local background may **not** be suitable for using the nods or alternate slit data for correction of zodiacal light if, for example, the area is crowded or contains extended emission around the science target.

Experience with the current IRS data suggests that it is not possible to effectively subtract the background from a single SH or LH pointing, even if the target is a point source, due to the small sizes of the slits and PSFs. In this case, observers should take an off-source (i.e., blank sky) pointing, or use data from a low-resolution module to set the continuum level. For observations of faint sources using the LH module, it is especially crucial that observers take off-source observations, close in time, in order to reduce the impact of rogue pixels. Observers should keep in mind, however, that off-source pointings can be costly in terms of total AOR duration with SH or LH observations, especially for single, deep exposures. On the other hand, if a spectral map with many pointings is being used, then a series of off-source pointings would add only a small overhead to the overall cost of the observations, and could be subtracted from all the

positions in the map.

Spot now provides an option to create automatically background AORs for SH and LH staring and scan map observations, by creating a new AOR with the same SH or LH integration time as the target observations. You will need to specify only the coordinates of the blank sky where you would like the background observations to be taken. Spot provides an input dialog where you can specify either an absolute position or an offset relative to the absolute target position for the background observation.

In addition, bright sources located even several arcseconds outside of the IRS slits can scatter light into the slit, which will contribute to the spectral background. In this case, off-source background observations that reproduce the scattered light conditions without having the science target in the slit will be important to subtracting the scattered light background component.

Since pipeline version S13, sky subtraction by nod difference in the low-res modules is performed and included in the products delivered to the observer. Observers can do their own sky subtraction for low-resolution data using either the two nod positions or the other sub-slit (if both SL1 and SL2 or LL1 and LL2 are selected). Alternatively, observers can subtract a model background from their data that corresponds to the precise date of observation. These constant background values are available via the background estimation routine in Spot. The corresponding calculation algorithm is discussed in the "Observatory and Instruments" section of the SSC website. Finally, the observer should consider planning dedicated, off-source background observations as part of their AORs if the methods described above are not expected to produce satisfactory results. After subtracting the sky, the observer can use the Spitzer IRS Custom Extractor (SPICE) software, which is available on the SSC website, to re-extract the 1D science target spectra.

In general, observers should plan off-source background spectra into their AORs if any of the following issues are relevant: (1) observations in high background or crowded regions, (2) presence of extended emission around the science target, (3) attempts to detect faint features or need for accurate continuum shape and/or absolute flux scaling, (4) use of either high resolution module (SH or LH), and (5) when scattered light is expected to be significant (for bright sources near the science target).

7.2.6 Best Observing Practices

In this section we present a summary of recommended practices to follow when designing an IRS observation. The list is neither exhaustive nor applicable to every science situation, and **reading this section should not be a substitute for a careful reading of the rest of the manual**. In particular, we recommend that observers pay careful attention to all material that is presented in **bold font** or (especially) material that is **enclosed in text boxes**. The latter usually

summarizes in one or two sentences an important aspect of IRS observation planning that is described at length in the surrounding text (see index under "IRS: boxed item" for a list of page numbers).

The basic characteristics of the spectrograph are summarized in the IRS Pocket Guide, available from the SSC website. They are also available in Table 7.1, Table 7.2, Table 7.5, and Figure 7.6. Assessing this information in light of your science goals should be the first step in determining whether to propose to use the IRS.

- *Peak-Up Accuracy*: Observers should be mindful of their required photometric precision, which will translate into peak-up accuracy. In general, high-accuracy peak-ups are not much more expensive than others and produce better results (spectroscopic repeatability of $1\sigma = 2$ to 5%), but medium accuracy peak-ups are appropriate for working with the wide slits (LL and LH). Low accuracy peak-up should not be used. See section 7.2.4.

- *Peak-Up Acquisition Targets*: The observer should ensure that he/she has picked up an isolated point source (no objects within $60'' \times 80''$) within $30'$ of the target by checking the MSX, 2MASS, and/or DSS images. The source should have a flux density within 5 mJy and 150 mJy in the blue or 15 mJy and 340 mJy in the red, with the exact minimum value depending on the local background level. This value should be entered correctly in Spot, as the peak-up source will be rejected if the observed flux is discrepant from the predicted flux by a factor larger than 5. The observer should take care to enter the epoch (the year for which the position was obtained), not to be confused with the equinox (which determines the actual coordinate system used) of the coordinates and the stellar proper motion. The later should be given in arcsec/year. Note that peak-up stars from the 2MASS catalog do not include proper motion information. See section 7.2.4.4.

- *Redundancy*: The observer should ensure that he/she has at least two cycles for each observation – this will provide greater redundancy for cosmic ray hits and rogue pixel identification. For example, rather than having 1 cycle $\times$ 480s, the observer should choose 4 $\times$ 120s.

- *Spectroscopic Saturation*: A-to-D saturation limits for the low resolution modules are ~5 Jy (10 Jy) at 10 μm (25 μm). For the high resolution modules, the limits are 60 Jy (70 Jy) at 12 μm (25 μm). This means that in the shortest possible observation (6 sec), at least one of the samples will be saturated. Given the characteristics of the SSC pipeline, saturation is not necessarily fatal although it will result in less signal-to-noise than expected. The observer should decide whether or not the loss of signal-to-noise in certain regions of the spectrum is compensated by a better

exposure elsewhere. See section 7.2.1.2.

- ***Imaging Saturation***: for the shortest exposure times, in low background, point sources brighter than 0.19 Jy at 16 μm will saturate the blue peak-up image. Point sources brighter than 0.42 Jy at 22 μm will saturate the red one. Extended sources brighter than 450 MJy/sr for both the blue and red will saturate the detector in the shortest available ramps. As in spectroscopy, source saturation is not necessarily fatal, as the image flux can usually be recovered. See section 7.2.3.4.5.1.

- ***Sensitivity***: In order to estimate the feasibility of his/her observations, the observer should use the sensitivity calculator available on the SSC website. Observers should be aware, however, that this calculator does not include any noise introduced by the extraction. It assumes perfect extractions with perfectly centered sources unaffected by rogues or fringing. In practice, the maximum continuum signal-to-noise ratio achievable in the low resolution modules is ~100, while in the high resolution modules it is ~50. See section 7.2.2.

- ***Peak-up array saturation and SL observations***: The current pipeline has difficulty performing the proper droop correction for SL data when the peak-up arrays are strongly saturated. All observers using the SL module should make sure that, for their observing window, the peak-up imaging fields do not fall in regions with very bright sources or backgrounds. The limits to avoid depend on the ramp times. At both 16 and 22 μm the limits are: 500 MJy/sr at 6 sec, 210 MJy/sr at 14 sec, 50 MJy/sr at 60 sec, and 25 MJy/sr at 240 sec ramp times. Observers should also avoid detector orientations that would put a source brighter than 41 Jy in any of the peak-up arrays, for ramps 240 sec long. As mitigation, users may constrain the time in which the observation will take place or reduce the ramp times (while increasing the number of cycles). See section 7.2.5.1.

- ***LL Observations of very faint sources:*** Observers interested in obtaining, in particular, LL spectra of very faint sources (~0.25 mJy at 24 μm) should consult the document, "Obtaining Spectra of Very Faint Sources with IRS," available at http://ssc.spitzer.caltech.edu/irs/imptnotes.html. In summary, the document recommends using spectral mapping mode (instead of Staring mode), moving the target along the slit between exposures, and dividing the total observations in multiple AORs (to get around the 8-hour limit for a single AOR). Furthermore, in order to minimize the amount of latent charge, the observer should request, via the AOR comment field, that the observation be scheduled immediately after "skydark" calibrations. This request should be described and justified in the proposal.

- ***Sky Observations, high resolution***: For the LH module, the SSC strongly

recommends that the observer plan for sky background observations close in time to the primary observation for sources fainter than 2 Jy at 25 μm. Sky subtraction will significantly reduce the impact of the rogue pixels in the signal-to-noise. For the SH module, sky subtraction is recommended if accurate flux calibration is needed. The sky observations for the high resolution modules should be observations of nearby empty fields, as extracting sky measurements from within the target observations is almost impossible with the high-res slits. See section 7.2.5.2.

- *Sky Observations, low resolution*: The sky can be subtracted using the two nod positions or observations in which the target is in the other slit in the detector. The latter are possible because the observer automatically gets 'off-slit' background measurements for (for example) order 2 when observing in order 1. If using off-slit sky observations, the observer should make sure that the off-slit regions are not contaminated by a bright source or extended nebulosity in the field of view. Also, the off-slit sky is only useful if the two orders have the same exposure time. Off-slit sky measurements may also be less accurate than nod-sky measurements, if separated by long times or if there is a large gradient in sky background. Observers interested in extended sources will need to assess whether or not these procedures are acceptable for their goals. If not, dedicated sky observations of empty fields will likely be necessary. See section 7.2.5.2.

- *Spectrophotometric Calibration*: The SSC provides the observer with photometrically calibrated products. The absolute photometric accuracy of point source spectra is limited by the accuracy of the models used for calibration and currently stands at ~10%. Representative slope deviations from the beginning to the end of an order with respect to point source models are 7% for the high resolution modules and 3% for the low resolution ones. Residual features in the flats produce systematic residual artifacts of 2 to 3%, which may affect detection of faint spectral lines. See section 7.3.3.2.

- *High-Res resolution:* Unresolved spectral lines are only Nyquist-sampled at the longest wavelength ranges in each module. If using the high resolution modules, the observer should understand that this will result in errors in the measured line widths. Observers should be careful when interpreting lines with $R \geq 400$ as resolved. See section 7.1.6. Observations of extended sources present unique challenges for slit spectrographs, since extended source astrophysical references of known surface brightness as a function of wavelength are not generally available. The SSC provides limited support for the spectrophotometric calibration of these kinds of observations in the form of aperture loss correction functions and slit loss correction functions, available via SPICE. The users interested in accurate calibration should make sure they have access to photometric observations (e.g., IRAS, MSX, ISO, MIPS, and/or the IRS

PUI arrays) in the appropriate wavelength range.

- *Mapping extended sources:* The IRS instrument offers a powerful form of spectral imaging, known as "Spectral Mapping mode," in which the telescope moves by small steps through a grid of positions, stopping at each step to collect one or more spectral exposures. Large regions can be mapped to relatively faint surface brightnesses, providing 3D spectral cubes covering extended objects. Information on these kinds of observations is provided at http://ssc.spitzer.caltech.edu/irs/documents/specmap_bop/ and in Section 7.2.3.3 of this document.

- *Detecting faint emission next to bright targets:* The PSF detected by the spectrograph has considerable structure, whose spatial scale depends on the wavelength of observation. Users interested in detecting faint emission next to bright targets should model the observed PSF using STINYTIM (at http://ssc.spitzer.caltech.edu/archanaly/contributed/browse.html). Very roughly, the PSF flux in the wings will be smaller than 1/100 of the peak source flux at distances larger than 1.5 to 2 slit widths from the peak (for the low-res modules). This estimate is provided as a general guideline only and should not replace a detailed PSF analysis using STINYTIM. For example, experiments with the simulation software STINYTIM indicate that, for a source with a brightness of 1 Jy at 8 μm, the extent of the PSF will be detectable at the 1 mJy level up to 10" away from the source. For a source that is 1 Jy at 25 μm, the 1 mJy level is reached at distances of ~30". Observers that are interested in detecting faint structures next to bright sources should take into account the shape of Spitzer's PSF when designing their AORs. It is strongly recommended that such target observations have dedicated sky observations at the same distance and orientation from known point sources, so that the diffraction pattern can be removed.

7.2.7 Astronomical Observation Template (AOT) Examples

Although a full description of the IRS AOTs is given in section 7.2.3, it is instructive to step through a few worked examples as a guide. See the Spitzer Observation Planning Cookbook (which is available on the SSC website) for additional detailed examples of IRS AOTs. (Note: examples were created using Spot version 15.0 and AOR estimator version 15.0.0. Future versions of the Spot software may produce slightly different numerical values from those shown here.)

7.2.7.1 Example #1: High-Resolution Staring Mode Observation of a Single Fixed Target

In this example, we will obtain a high-resolution spectrum of a bright target, in order to search for two faint emission lines at 15 and 26 μm. We will use both the SH and LH modules. The target is at RA (J2000) = 15:34:57.4, Dec (J2000)

=+23:29:52, and has IRAS 12 and 25 µm flux densities of 0.5 and 7.9 Jy, respectively. By placing this object in the Spot target list, we can check its visibility, showing that there are many windows of opportunity for Spitzer observations (see Figure 7.41). Recall that the orientation (roll angle) of Spitzer is **not** observer-selectable, and depends on the date of observation. For example, if this target is observed on 01 April 2006, then the SH and LH slits will have positions angles on the sky of 121.4º and 36.6º East of North, respectively. Note that a particular observation date should not be specified unless scientifically essential, because of the resultant constraints imposed on scheduling the observation. This feature of Spot is most useful to simply determine the range of dates during which a particular object could be observed.
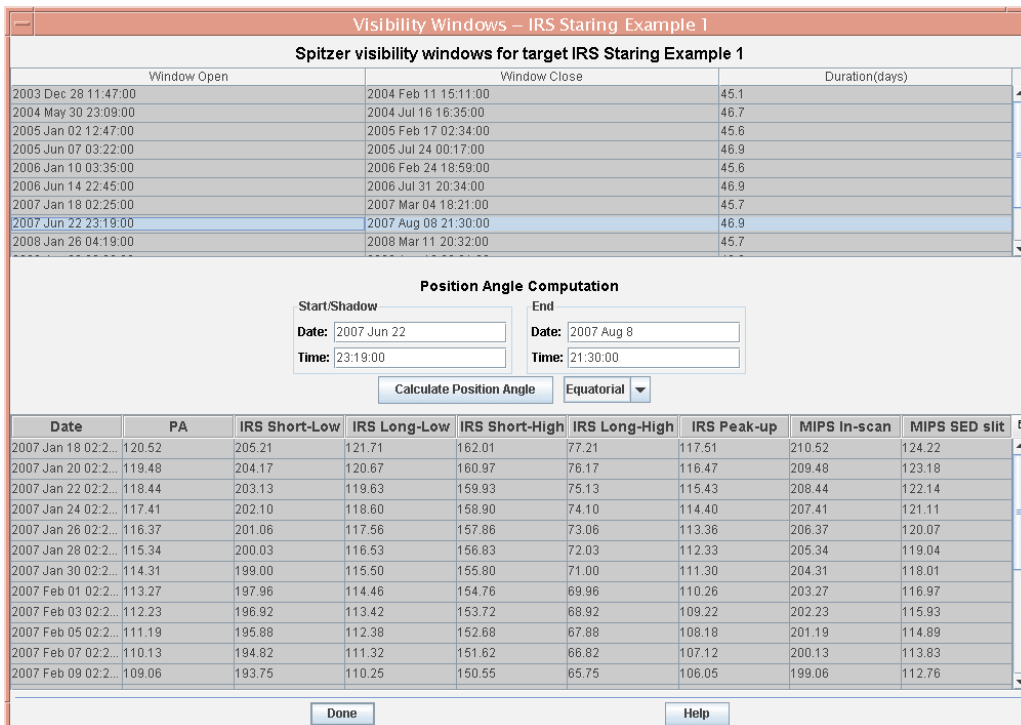
Figure 7.41: Target visibility window for Example #1.

We would like to obtain at least a 10σ detection (or limit) on the two emission lines, which are expected to have fluxes of approximately 5×10^{-18} W m^{-2}. This target is located at an ecliptic latitude of approximately 41 degrees, corresponding to the MEDIUM background case. The appropriate sensitivity curves (see Figure 7.21 and Figure 7.22) indicate that a single, 120 second exposure will provide a signal-to-noise of unity in the SH module for an unresolved emission line having a flux of approximately 6×10^{-19} W m^{-2} at 15 µm. Similarly, a 240 sec exposure will provide a signal-to-noise of unity in the LH module for an unresolved emission line having a flux of approximately 6×10^{-19} W m^{-2} at 26 µm. If we select two cycles per observation, and the 120 second and 240 second exposure times for the SH and LH slits, respectively, we should obtain a signal-to-noise of approximately 16–17 in the final spectra from both modules for our target line

flux of 5×10^{-18} W m^{-2}. Note that since the default in staring mode is to obtain two spectra per slit, by selecting two cycles per slit we will actually obtain four 120 second exposures in the SH module, and four 240 second exposures in the LH module (combining all four spectra is assumed in the sample S/N value given above). Also, while the target is moderately bright (but below the BSL for SH and LH), there is no problem with saturation in the high-resolution modules, as the flux density in the observable wavelength regime is well below 70 Jy (see Figure 7.23). The observation is shown in the completed AOT in Figure 7.42.

Figure 7.42: Completed IRS Staring mode AOT for Example #1.

While the science target is a point source, it is too bright in both IRS peak-up filters to serve as the peak-up target (based on its known IRAS flux densities). Therefore, we have selected a fainter, nearby star using the "2MASS selection" routine, and entered this into the AOT for a High accuracy peak-up (see Figure 7.42). Note that we must locate the proper motion values for the selected peak-up star from another source (e.g., the Tycho-2 catalog) as this information – which is crucial to the success of the peak-up – is **not** available from 2MASS. Once the

AOT is filled out, the estimated AOR duration should be computed by clicking on the "Calc. Obs. Time" button at the bottom of the form, the results of which are shown in Figure 7.43. The total AOR duration in this example is less than the 8-hour limit for IRS AORs.

Figure 7.43: Resource Estimates for Example #1.

7.2.7.2 Example #2: High-Resolution Observations of a Fixed Cluster Target

In this example, we obtain high-resolution spectra using both the SH and LH modules on a target (NGC 6946) located at RA (J2000) = 20:34:52.34 and Dec (J2000) = +60:06:14.12. In addition, we obtain spectra at several positions forming a cross-like pattern along the cardinal axes at offsets of 180″ and 360″ from the target coordinates. This is achieved by defining the target type as "Fixed Cluster-offset," as shown in Figure 7.44. Note that the coordinates of the center of the cluster are specified at the top of the target form, and the offset positions are given in the table. East and north offsets (in arcseconds) are given as positive numbers, while west and south offsets are given as negative numbers.

The completed AOT for this example is shown in Figure 7.45. We requested a Moderate accuracy blue peak-up on a nearby 2MASS star, with 1 cycle of 30 sec integration for the SH module and 1 cycle of 60 sec integration for the LH module. The default behavior in Staring mode is to take two spectra per slit so we will actually obtain two 30 sec exposures in SH and two 60 sec exposures in LH. The Resource Estimates window (see Figure 7.46) shows that the total AOR

duration for the nine positions is less than the 8-hour limit for IRS AORs.

Figure 7.44: Fixed Cluster-offsets target form for Example #2.

Even though this is an efficient way to obtain spectra for nearby locations without being charged for extra slew overheads, the observer should keep in mind two important restrictions. The first is that all offset positions must lie within a circle of 1° radius. The second is that the observing parameters (modules/slits used, duration, peak-up accuracy, and slit position angles) will be the same for every position in the cluster.

7.2.7.3 Example #3: Spectral Mapping for Observations Between 7.4 and 14.5 μm

In this example, we will use the SL IRS module to take spectra of an extended region using Spectral Mapping mode. When planning low-resolution observations, Spot considers each order of the slits separately. Assume for the moment that we are only interested in the spectral range between 7.4 and 14.5 μm covered by the first order of the SL module (SL1). In the example shown in Figure 7.47, we use the "Low-Res 7.4–14.5" (i.e., SL1) sub-slit to make an 8×2 grid of 14 sec exposures on NGC 1068 in the following manner: after the first 14 sec exposure (Position 1), we move the slit in a direction perpendicular to the slit length (white arrow) by 10″ and take a second 14 sec exposure (Position 2). We

repeat this procedure for a total of 7 times with a fixed step-size of 10″ so that in the end (Position 8) we have observed at 8 different positions along a direction perpendicular to the slit. The last pointing of the slit is 70″ away from the first position. We then displace the slit 50″ in a direction parallel to its long axis, and perform a set of exposures in reverse order from the first eight, but at the new displaced positions (Positions 9, 10, …, 16). It can be seen from the figure that by the time the final exposure is made (Position 16), the slit has returned to its original position in the perpendicular direction, but is displaced from the starting point (Position 1) by 50″ in the parallel direction. The length of the SL 1st order sub-slit is 57″ and its width is 3.7″, so in this example we have obtained spectra (in the requested sub-slit) over a region of 107″ by 70″.

Figure 7.45: Completed IRS Staring mode AOT for Example #2.

The completed AOT form for the map grid of Example #3 (illustrated in Figure 7.47) is shown in Figure 7.48. We have disabled peak-up because the exact

pointing for this large map is less important than if we were obtaining a single (Staring mode) spectrum across the target. The total duration time for this observation is shown in Figure 7.49. In Spectral Mapping mode, unlike Staring mode, we obtain only one spectrum per slit position, and so the total on-source exposure time for this sequence is 16×14 sec = 224 sec.
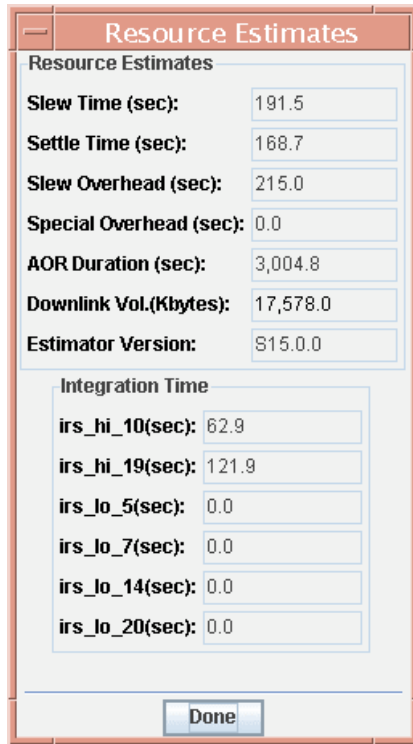
Figure 7.46: Resource Estimates for Example #2.

7.2.7.4 Example #4: High-Resolution Observation of a Moving Target

In this example, we obtain a high resolution spectrum using the SH module for a moving target, Europa, which is a bright moon of Jupiter. We are interested in observing it with a minimum of stray light from Jupiter. To do this, we will apply timing constraints that specify that we want to observe Europa at or close to the time of greatest elongation from Jupiter, as viewed by Spitzer. We determine the time of greatest elongation using the JPL Horizons tool with Spitzer as the coordinate center (see the Proposal Kit section of the SSC web site for detailed instructions). This tool returns the information that Europa will be within about ±2 degrees of maximum elongation between 05:20–06:50 on 01 March 2006.

The planned AOR is about 20 minutes in duration (see Figure 7.52). Consequently, we specify the timing constraints in Spot so that the AOR will start execution between 05:20 and 06:30 on 01 March 2006 (see Figure 7.50). This specified window is 20 minutes shorter than the window that bounds the target within ±2 degrees of maximum elongation. This was done because the timing

window specifies the interval in which the AOR can *start*, so in this case, truncating the window by 20 minutes (the duration of the AOR) ensures that Europa is observed while it is still within ±2 degrees of maximum elongation, even if the observation is scheduled to start at the very end of the specified timing window. The completed AOR for this example is shown in Figure 7.51. We requested no peak-up because the target itself is much too bright to be used for this purpose, and stationary offset targets (i.e., stars) cannot currently be used to peak-up on moving targets. Because of Europa's brightness, which is close to the saturation limit, we selected the shortest possible integration time available for the SH module. The total AOR duration is summarized in Figure 7.52.
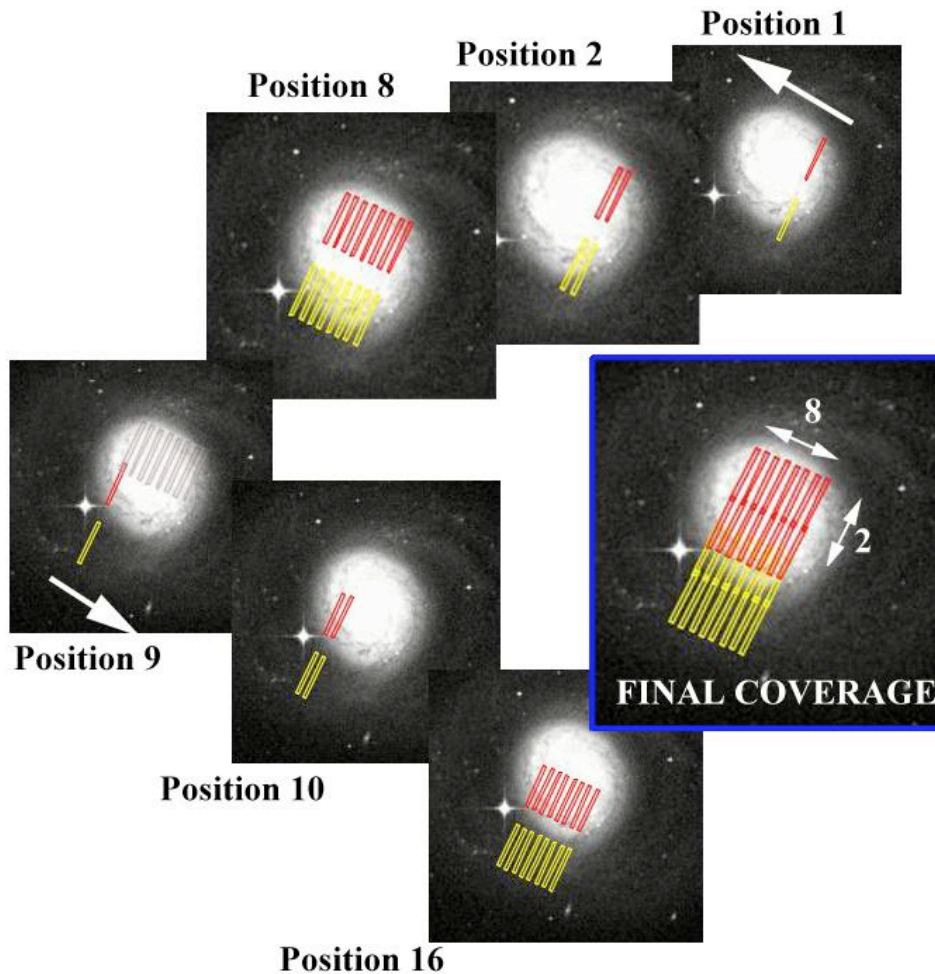
Figure 7.47: A schematic diagram of the observing sequence of NGC1068 for Example #3 on a particular observing date (hence, a particular slit position angle) taken from the animation sequence available in Spot. Note that although a given AOT refers to a specific sub-slit (dark or red), the other sub-slit (pale or yellow) also obtains data simultaneously. In certain circumstances, such as large fields, the fact that data are obtained in both sub-slits simultaneously can improve mapping efficiency if spectra from both sub-slits are required.

7.2.7.5 Example #5: Peak-Up Imaging Observation

This example demonstrates the basic procedure for using PUI to map a region

previously mapped with IRAC or MIPS. The field-of-view of the peak-up arrays is considerably smaller than the 5×5 arcminutes covered in a single IRAC or MIPS 24-µm exposure. Nonetheless, observers may choose to perform a 16 µm survey of targets previously observed with IRAC and MIPS, in order to fill in the wavelength gap between the latter instruments. Consider the case of mapping the galaxy NGC 7552 with the blue peak-up imaging array. This galaxy is comfortably observed in a single IRAC or MIPS pointing but requires multiple IRS peak-up array pointings for full spatial coverage (see Figure 7.53). To avoid gaps in the coverage, we want to use PUI mapping in array coordinates and choose steps that are nearly the size of the array. In a 6×4 grid, we can cover most of the required field-of-view. The change in rotation angle between observing dates makes it difficult to match the coverage exactly without substantially more map positions. In order to mitigate bad pixels, we use two positions in the cycling dither pattern.
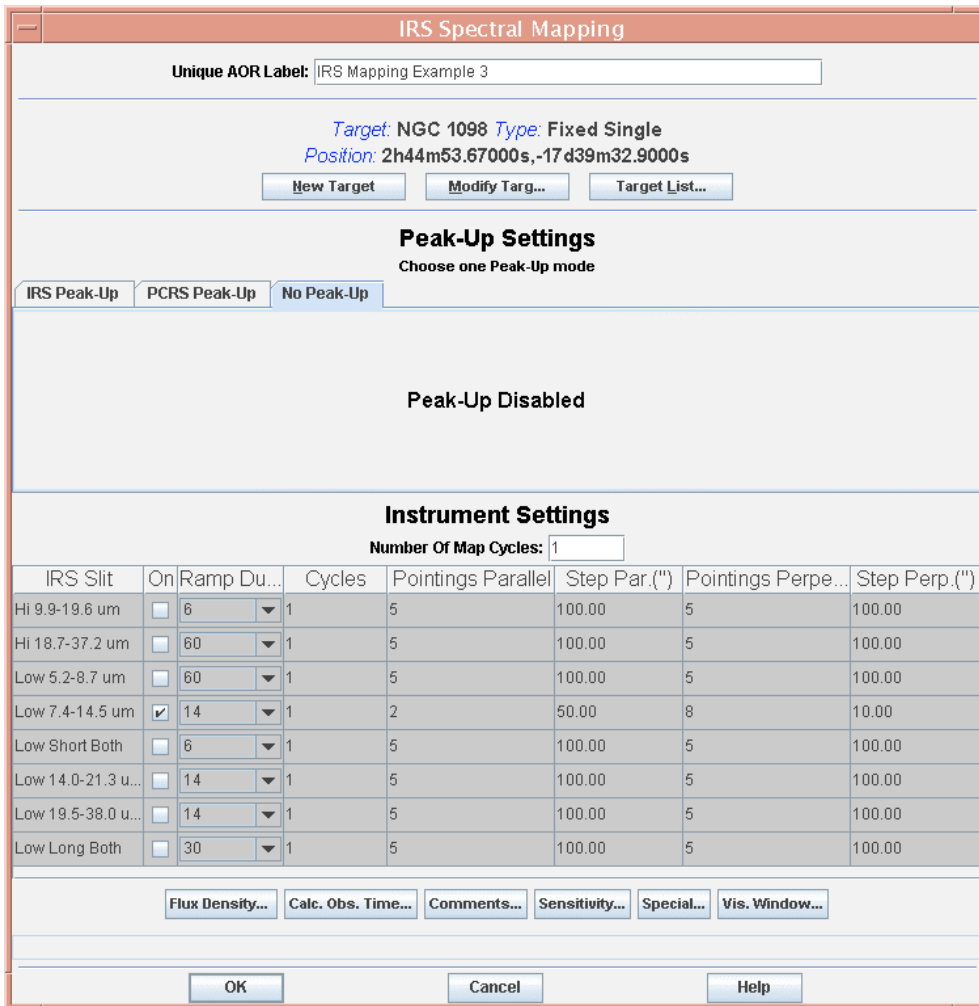
Figure 7.48: Completed Spectral Mapping mode AOT for Example #3.

Another consideration when planning PUI observations is the measurement of the local sky. In this example, we may choose to use the outlying parts of the grid to sample the sky and make additional measurements. However, we might instead add more grid positions to ensure that the outer regions of the target galaxy do not affect the background measurements. Alternately, we could choose to repeat the entire map at a different location. The completed AOT for this example is shown in Figure 7.54.

Figure 7.49: Resource Estimates for Example #3.

7.3 Data

An up-to-date discussion of the IRS data is available in the Data Handbook. Unresolved issues regarding data quality are discussed at http://ssc.spitzer.caltech.edu/irs/features.html. Here we provide a summary view of the more detailed discussion from the Data Handbook.

7.3.1 Instrument Calibration

The IRS is fully calibrated and produces spectrophotometric and photometric data. The wavelength scales, spectral resolutions, photometric responses, and noise characteristics of the four modules were initially determined through a series of laboratory measurements, and have been verified and updated on orbit. The spatial resolution, system response, and color corrections of the peak-up arrays are determined in a likewise manner. Table 7.10 lists the set of calibration measurements performed on the ground and scheduled for flight operations. Updates will be provided in future versions of this manual, or on the SSC website, as they become available.

Specific IRS calibrations that are performed in orbit include:

AOR Timing

Timing Windows (UT)

AOR Name: IRS Moving Object Example 4

Start date	Start time	End date	End time
2007 Sep 5	05:20:00	2007 Sep 5	06:30:00

Visibility Windows

Window Open	Window Close	Duration(days)
2003 Dec 05 17:05:00	2004 Jan 14 00:22:00	39.3
2004 May 03 21:16:00	2004 Jun 12 17:42:00	39.9
2005 Jan 11 03:10:00	2005 Feb 19 04:47:00	39.1
2005 Jun 11 07:20:00	2005 Jul 21 23:50:00	40.7
2006 Feb 21 16:48:00	2006 Mar 29 19:38:00	36.1
2006 Jul 21 03:14:00	2006 Aug 31 07:15:00	41.2
2007 Mar 31 20:48:00	2007 May 08 23:18:00	38.1
2007 Aug 31 15:49:00	2007 Oct 11 22:16:00	41.3
2008 May 10 06:16:00	2008 Jun 19 10:53:00	40.2
2008 Oct 13 04:42:00	2008 Nov 23 12:52:00	41.3
2009 Jun 23 17:03:00	2009 Aug 03 14:21:00	40.9

Visibility windows shortened to avoid: earth, moon

OK	Cancel	Help

Figure 7.50: Setting narrow timing constraints for Example #4.

Spectrophotometric Calibration: Observations of A dwarfs and K giants are used to establish the absolute flux calibration of the IRS and to update the flat fields for the high-resolution modules. These sources are located primarily in the northern and southern CVZs, but some sources (particularly A dwarfs) are located in the Operational Pointing Zone (OPZ) of Spitzer. Absolute photometric calibration of the peak-up sub-arrays relies on the fainter stellar calibrators, with overlap in the sources to be observed for the low-resolution spectral flats. The calibration is based on synthetic spectra of the standards as well as on templates provided by combinations of observed and synthetic spectra. Calibrations involving the zodiacal light rely on a model based on COBE/DIRBE observations.

Figure 7.51: Completed Staring mode AOTfor Example #4.

***Flat fields*:** Flat fielding for the low-resolution modules relies on a combination of observations of the zodiacal emission at low ecliptic latitudes and stellar calibrators. The high-resolution modules are not sufficiently sensitive to make use of the zodiacal light for flat fielding. Flat field calibrations of the peak-up arrays similarly rely on dithered observations of sky positions with maximum zodiacal emission. Pixel-to-pixel sensitivity is tracked in time through the use of internal stimulators.

Wavelength calibration: Wavelength calibration was initially established on the ground using a monochromator, for the purposes of verifying design requirements. The monochromator measurements provided 5 points per low

resolution order and 9 points for each of the high resolution modules, to give coarse calibrations with uncertainties of up to 0.5 μm. In orbit, the wavelength calibration has been fully rederived, by observing luminous Be stars and planetary or reflection nebulae, giving up to 25 points for the high resolution modules. Unresolved line centroiding accuracies of 1/5 resolution element (0.003–0.011 μm for SH and LH, and 0.015–0.11 μm for SL and LL) are achievable from flight data.

Figure 7.52: Resource estimates for Example #4.

Spectral resolution and line profile characterization: Accurate, in orbit measurements of the spectral resolution and instrumental line profile are obtained as a by-product of the wavelength calibration observations of PNe. Since the nominal spectral resolution of the SH and LH modules is around 600 and the expected line widths of PNe are less than 30 km/s, most lines will be effectively unresolved and their line centroids and line spacings within each order can be used to refine the pre-launch wavelength calibration of the modules. Unresolved lines are also used to accurately determine the shape of the instrumental line profiles.

Linearity: Non-linearity in the photometric response has been characterized on the ground and was further refined during IOC/SV. For the low-resolution modules this is achieved by observing the zodiacal emission at low and high ecliptic latitudes. For the high-resolution modules both external and internal (stimulator) sources are used to measure linearity.

Beam profile measurement: Point sources, such as stars and distant, luminous

infrared galaxies, are observed to characterize the point-spread function at various wavelengths.
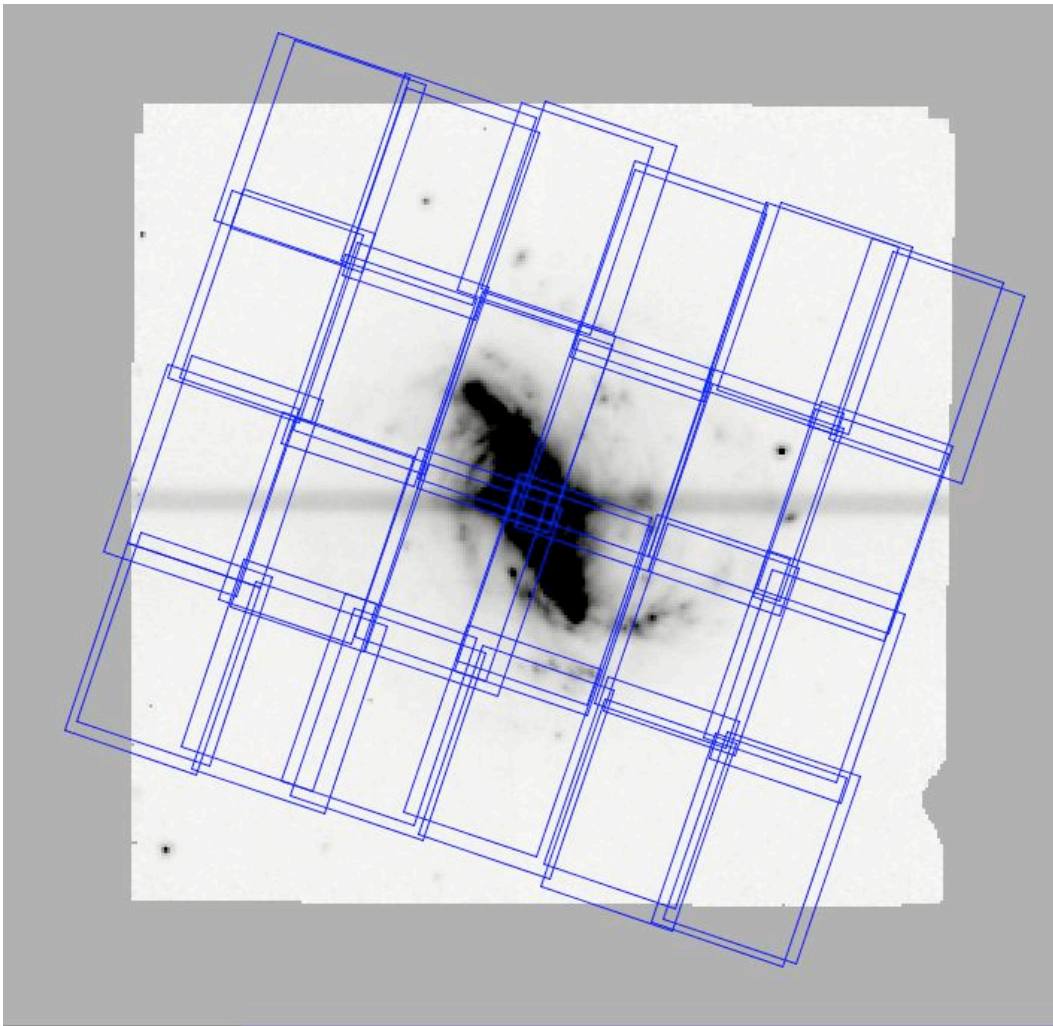
Figure 7.53: Spot visualization overlay for PUI mapping of NGC 7552. The underlying image is the IRAC image observed by the SINGS Legacy project retrieved from the Spitzer archive.

***Spectral leaks*:** An assessment of out-of-band spectral leaks in the IRS can be performed by observing known sources which have significant IR emission outside the bands of interest. Suitable objects for investigating blue leaks are luminous infrared galaxies with already known, relatively featureless mid-infrared spectra, as well as asteroids and some planets (e.g., Uranus and Neptune), relying on reference models of all of these objects. While some asteroids may be expected to vary in intensity over their rotation period due to surface albedo variations, the rotation periods of the candidate asteroids are long compared to the necessary IRS integration times. Departures from a thermophysical model (in which the shape, cratering, thermal conductivities, and periods may be assumed) should be of a broadband nature that can be distinguished from a blue light contamination. Neptune can also be observed, but is not always in the OPZ of

Spitzer and is too bright for the LL module (1st order).

Figure 7.54: Completed PUI AOT form for Example #5.

Droop characterization: Droop is a change in the measured charge on a given pixel, and its magnitude is a function of the total flux falling on the array. Ground tests suggest that this effect can be modeled globally for each array, by using a single coupling coefficient. On the ground, these droop coupling constants have been estimated by analyzing data from masked regions on each array. During IOC, the model was tested using the unilluminated areas of each array. Since "droop" appears only in observations acquired in RAW data mode, a direct measurement/evaluation of the droop is achieved by comparing calibration data obtained in both SUR and DCS modes.

Scattered light: The presence of scattered light and/or ghosts on the array has been tested on the ground and limits of their contribution have been established by observing bright sources. The same technique was followed during IOC/SV. An added test to examine the effects of scattered light from the telescope baffle

was performed during the cool-down of the instrument in the first weeks of IOC. Since light from the temporarily warm baffle then dominated any sky emission, it was distinguishable because its signal followed the telescope cool-down curve.

***Slit positions*:** The mapping of the focal plane to obtain precise positions and orientations of the slits with respect to the IRS Peak-Up arrays and the PCRS was carried out in the course of a multi-step campaign during IOC by scanning a star numerous times across each slit. The PCRS is used to define the telescope boresight, and then the positions of the IRS Peak-Up arrays and the individual slits are measured separately with respect to the PCRS to an ultimate accuracy of 0.14″ (radial) for the short wavelength slits, and 0.28″ (radial) for the long wavelength slits. Note that this accuracy is considerably higher than the High accuracy requirement of 0.4″, which will limit the absolute photometric uncertainties to between 2 and 20%, depending on wavelength (see Table 7.10).

***Dark current measurement*:** As the IRS does not have a shutter, dark current is measured from dithered observations on the darkest patches of the sky, as part of every observing campaign. These observations effectively provide a measurement of dark current plus sky and telescope emissivity contributions. The frequency of the dark current measurements depends on stability, as evaluated during IOC and monitored during the routine phase of operations, and may be adapted to changes over the course of the mission and specific needs of the campaign. Integration times will match those utilized during the campaign. The north ecliptic pole (always in the northern CVZ) will serve as the fiducial dark patch, but a number of measurements will also be made towards the Spitzer deep fields, which will be observed extensively by MIPS and IRAC. Variations with solar elongation in the low-level zodiacal emission at the north ecliptic pole are estimated to be extremely low, but will be monitored over the orbit.

Table 7.10: IRS calibration measurement description

Sequence	When	Duty Cycle	Purpose	Description
Filter Transmissions	Ground		Necessary for color correcting data.	Transmission vs. wavelength.
Detector Response	Ground		Necessary for color correcting data.	Responsivity vs. wavelength.
System Response	Ground		Color correction of peak-up cameras and checking spectral leaks.	System Response vs. wavelength.
Performance Verification	Ground		Verify basic detector performance and establish array characteristics.	η/β, ηG, dark current, and read noise, pixel offsets, etc.
Linearity	Ground, IOC		Establish operating range of detector and measure non-linearity.	Vary int. time with fixed flux. Vary source flux at fixed integration time.
Slit Position	Ground, IOC		Determines slit positions relative to peak-up camera.	Get central axis and orientation of slit.
S-Curve Distortion	IOC		Correct data for geometric distortion.	Place star at several positions along the slit.
Wavelength	Ground,		Determine wavelength	Observe planetary and

Sequence	When	Duty Cycle	Purpose	Description
Calibration	Normal Ops	Every 24 hours	calibration of spectrographs.	reflection nebulae, and bright Be stars with emission lines.
Spectral Resolution	Ground, Normal Ops	Each campaign	Determine resolution of spectrographs.	Observe extended source with narrow emission lines (PNe).
High-Low Zodiacal Flat field	Normal Ops	Each campaign	Determine flat field for faint sources.	Use difference at low and high ecliptic latitudes.
Dark measurement	IOC, Normal Ops	Every 12 hours	Determine dark performance of the detectors.	Observe darkest region in northern CVZ near North Ecliptic Pole at all integration times. Dither and repeat to remove cosmic ray effects and possible background sources. Compare with observations of MIPS deep fields when viewable.
Extended Source Flat field	Normal Ops	Each campaign	Determine flat field for bright sources.	Dither a bright-extended source along the slit.
Star Scan Flat field	Normal Ops	Each campaign	Check flat field for all sources.	Slew a star along the slit.
Pixel Offsets	Normal Ops	Each campaign	Get pixel offsets (electrical + dark current) for each integration time.	Use signal at high and low ecliptic latitudes from the high-low zodiacal flats.
Spectrophotometric Monitoring	Normal Ops	Every 4 hours for first 5 campaigns, then every 12 hours	Tie calibration to known well-behaved sources.	KIII and AV stars.
Spectral leaks	IOC		Check for spectral leaks in filters.	Observe point-like red sources (e.g. Mrk 270) and compare to blue sources.
Beam Profile	IOC, Normal Ops	Passive	Determine beam profile of cameras and spectrograph.	Observe point sources, such as stars and LIRGs.
Scattered light and Inter-order Cross-talk	Ground, IOC		Determine corrections for scattered light and optical cross-talk.	Place bright source(s) on and off the slits and peak-up cameras.
Array and System Stability	Normal Ops	Passive / Each campaign	Determine how often calibrations need to be performed.	Temporal variability of pixel offsets and response. Do stimulator flash every time a module is turned on and before it is turned off.

Notes to Table 7.10:
- Duty cycle refers to in-orbit measurements only.
- Observations done "Each campaign" were also done during IOC.
- IOC-only measurements will be checked or repeated infrequently after IOC.

7.3.2 Data Products

This section is a high-level overview of IRS data products. For more information,

please see the IRS Data Handbook (also see section 7.2.3.4 for a brief description of IRS PUI data products).

There are three basic IRS spectroscopic data products. The IRS Raw Data consists of one or more Data Collection Events (DCEs) that have been converted to FITS format. Each IRS DCE consists of a number of reads (samples) of the array at fixed time intervals without a reset between the samples. All IRS spectra will contain four, eight, or 16 reads of the array between resets. Thus, the raw FITS data will be a multi-plane data cube with each data plane representing a read of the array. Any ramp duration (exposure time) longer than 16 seconds still results in a sixteen plane raw FITS data product. The archive will release these data cubes under the name of "raw.fits" to the users.

The IRS *Basic Calibrated Data (BCD)* is a FITS file made from each DCE, which has all the "Spitzer-specific" effects removed. For the IRS spectra, the basic calibrated data are either (1) a calibrated, 2-D echellogram for the high-resolution data, or (2) a calibrated, long-slit 2-D spectrum for the low-resolution data. The BCD FITS files are accompanied by additional FITS files used in the processing (e.g., the wavelength map, flat field, efficiency correction, and traceable error images). A single DCE results in a single BCD. An IRS AOR can be made up of many DCEs.

The IRS *post-BCD product* is intended to convey the richness of the target spectrum and allow for a "first look" scientific analysis. The SSC will provide extracted, coadded 1-D spectra, with some measure of sky subtraction. Flux is assigned during the extraction process, using electrons/second/Jy conversions obtained from measurements of stellar standards. The BCDs are therefore extracted on the assumption that the calibrated signal originates from a point source, and the 1-D spectra will be in units of Jy vs. µm and provided to the users in ASCII table format.

The above discussion applies to IRS spectroscopic and peak-up imaging data only. Although no pipeline processing is applied, the IRS peak-up target acquisition data are also transferred to the ground and distributed to the observer. As described in section 7.2.4.1.6, the IRS peak-up target acquisition process uses the Double Correlated Sampling technique, and all data processing is handled on board. Each IRS Peak-Up produces six unprocessed DCS (difference) images, as well as two dark-subtracted, cosmic-ray cleaned, flat-fielded and background subtracted "average" frames. These average frames are the ones that are used by the algorithm to find the peak-up target. Therefore, the observer selecting an IRS Peak-Up will receive, in addition to their spectral data, a set of eight peak-up images.

It follows from this description that the Spitzer archive will not only include the Raw, BCD, and post-BCD images, but also additional information, such as uncertainty images and mask files. These products will allow observers to assess processing of the pixel data at certain junctures in the pipeline. Full reprocessing

of campaign data will occur at certain times during the mission, as needed by improvements in the processing algorithms or calibrations.
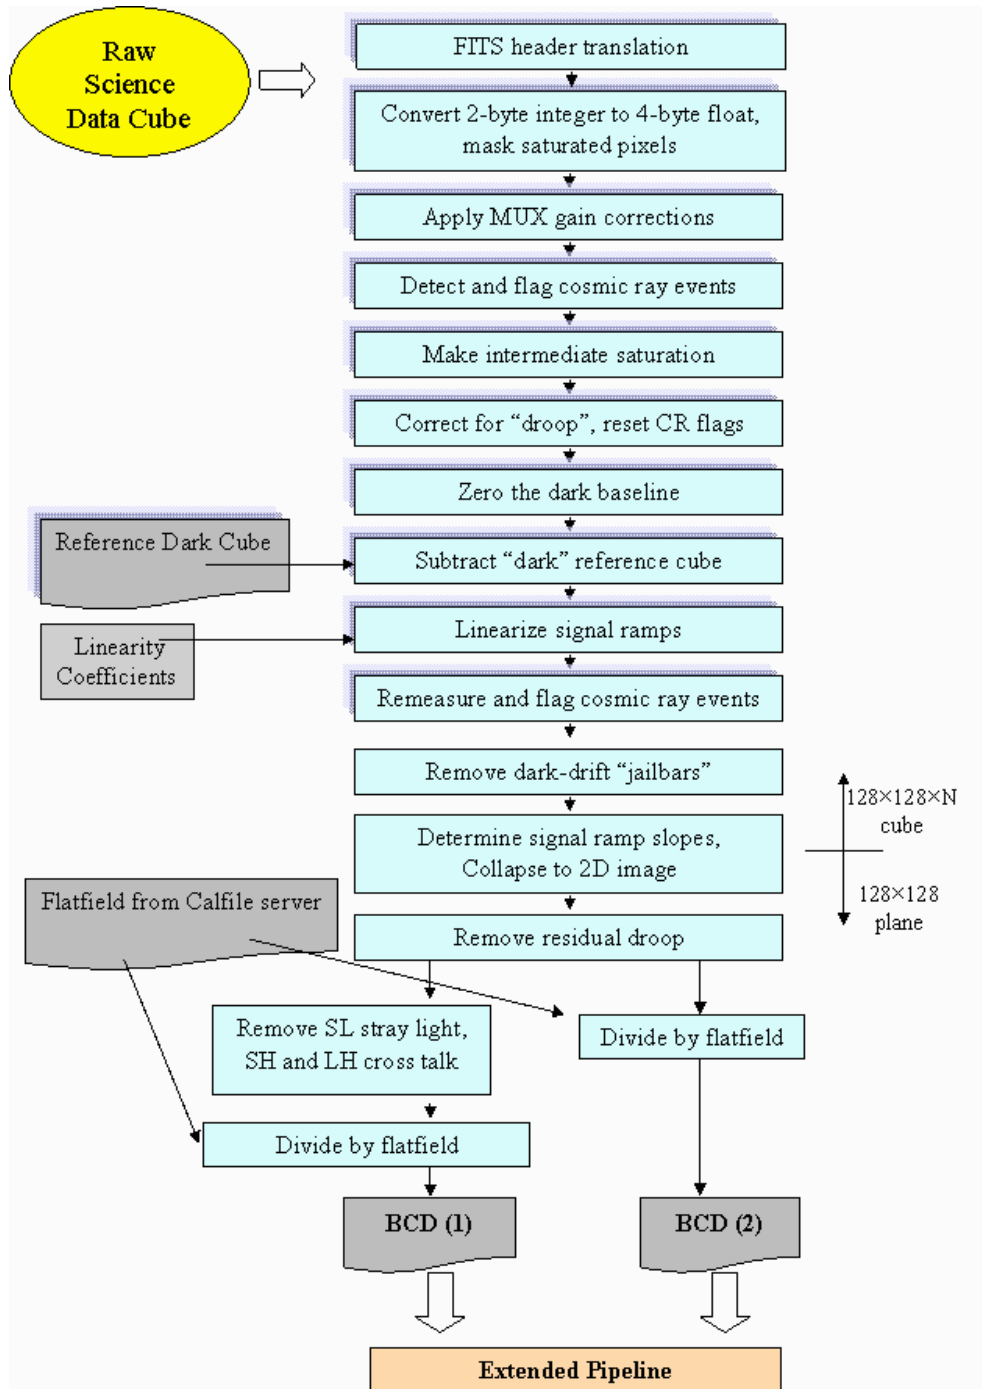
Figure 7.55: The IRS Science Data Pipeline. BCD (1) and BCD (2) represent planar images that either have been corrected for peak-up stray light or echelle order optical cross-talk, or remain uncorrected, respectively. "Shadowed" steps are meant to indicate processes that are performed on the SUR (128128N) data cubes, while the other steps are performed on the 128128 planar images. Calibration inputs to the science pipeline are indicated on the left. The extended pipeline is shown in Figure 7.56.
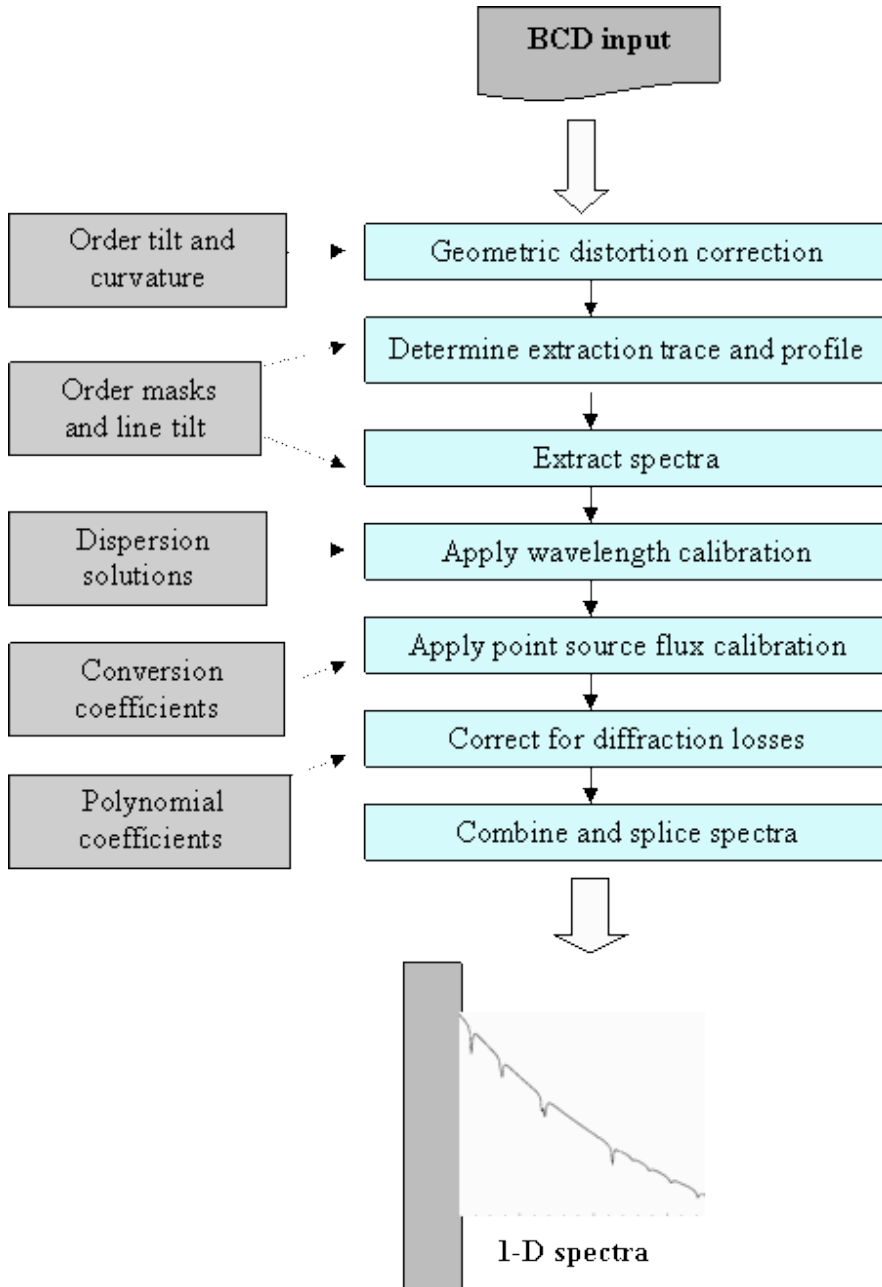
Figure 7.56: The IRS extended pipeline

7.3.3 Data Processing

The IRS data processing performed at the SSC occurs in two separate, but related, pipelines. In the Calibration Pipeline, data obtained from calibration observations are processed to produce files (data cubes or planes) and coefficients that will be used to remove instrumental signatures and perform spectral and photometric calibrations. Separate "threads" in the Calibration Pipeline are devised according to the task, and are generally run offline. The calibration files and coefficients generated from the different threads are validated and may be combined (if applicable), and stored for use in the Science Data Pipeline to process IRS science

observations. The Science Data Pipeline is divided into a BCD thread, for processing each raw science data cube (on a DCE basis) into a fully calibrated planar image, and into post-BCD threads, for extraction of BCDs into 1-D spectra, coaddition of BCDs, extraction of coadded BCDs, and coaddition of 1-D spectra. Rectified BCDs are also created in the post-BCD threads, but spectra are not extracted from them. The major components of the IRS pipelines are described as follows; also see Figure 7.55 and Figure 7.56.

7.3.3.1 Calibration Pipeline

The calibration pipelines are designed to pre-process observations obtained for dedicated calibration activities, from AORs using celestial objects, or IERs using celestial sources or internal measurements. Intermediate and final output of the calibration threads are further processed in interactive analysis (I/A) environments. Only the processing and application of dark current measurements to science observations is done automatically, although these are monitored offline as well. The set of calibration processing and offline I/A procedures perform the following tasks.

1. ***Order Mask Generation*** – Locates the position of each order in the SH and LH data, and determines the order widths, curvatures, and tilts, which are optimized according to wavelength solutions and line-tilt inputs.
2. ***Dark Current Determination*** – Processes and combines dark sky observations for subtracting the dark currents from each array. Applies the combined dark current measurement obtained nearest in time to the science observation, at matching integration times. Offline analysis monitors the accuracies and stabilities, and combines many dark measurements to create low noise "fallback" darks that are used with initial processing of science observations (see section 7.3.3.2).
3. ***Non-linearity Measurement*** – Calculates the photometric non-linearity coefficients for every pixel in each array.
4. ***Bad Pixel Mapping*** – Updates the bad pixel map for all arrays.
5. ***Flat Field Frame Generation*** – Generates a flat field frame for each array.
6. ***Dispersion Measurement*** – Calculates the wavelength solution and line-tilt for each array.
7. ***Slit-to-Pixel Mapping*** – Matches physical locations along each slit to position on the array, and provides geometric distortion corrections.
8. ***Short-Low Stray Light Correction*** – Measures and corrects for light falling on the SL spectral orders from the peak-up apertures.
9. ***Order Cross-Talk Correction*** – Measures and corrects the optical cross-talk between orders for the SH and LH arrays.
10. ***Fringing Measurement*** – Measures the fringe amplitudes, phases, and frequencies for the SH and LH modules.
11. ***Flux Calibration*** – Calculates the absolute flux conversion scale (electrons/sec to Jy) for each order.
12. ***Order Tuning*** – Determines polynomial coefficients to correct for point-source diffraction losses and higher order flat fielding residuals in the

photometric calibrations.

7.3.3.2 Science Data Pipeline

The science data pipeline performs basic conversions of the raw data values, organizes header information obtained from the uplink tables and telemetry, and applies the calibration files and coefficients to remove instrumental signatures, producing a BCD from each science DCE obtained with the SUR data mode. Spectra are extracted, wavelength-calibrated, and flux-calibrated in the Extended Science Data Pipeline. The tasks performed in the Science Data BCD and Extended pipelines are listed below. The flow is depicted in Figure 7.55 and Figure 7.56. It should be noted that the flow charts only represent the data processing flow. Several operations occur while the pipeline is running, such as data type and integrity sanity checking, measurement and database storage of quality analysis items, and FITS header updates.

BCD Pipeline Tasks (in order of occurrence in the pipeline):

1. ***Header translation*** – Filters and converts uplink information and telemetry accompanying each DCE into FITS keywords.
2. ***Data conversion and saturation detection*** – Converts raw data values in each pixel from unsigned integer to floating point type. Detects and masks digitally saturated pixels.
3. ***Gain correction*** – Corrects for the read-out channel-dependent gain differences in the arrays. Converts signal to electrons.
4. ***Cosmic ray identification*** – Identifies and flags cosmic ray events along the signal ramps.
5. ***Saturation correction*** – Makes intermediate corrections to saturated pixels for the purpose of droop correction, then resets the cosmic ray event flags.
6. ***Droop corrections*** – Corrects the signal on the arrays (on a pixel by pixel basis) for the coupling of signal between pixels (illumination "droop" effects).
7. ***Dark baseline drift correction*** – Re-establishes a dark current zero point for the purpose of dark current subtraction, by referencing unilluminated regions of the arrays.
8. ***Reference (dark current) subtraction*** – Removes a reference DC offset level from the science data. The operation is performed twice: first, when the science data are received at the SSC and initial processing is performed using "fallback" or default dark current cubes, then again during a full reprocessing of all data in the campaign when dark current measurements in the same campaign have been processed and the most contemporaneous dark cube of the appropriate integration time is subtracted. First and second processings of the data will be inspected by the IRS IST and DQA systems.
9. ***Linearization*** – Linearizes the science data signal ramps.
10. ***Cosmic ray correction*** – Remeasures and flags cosmic ray events along each signal ramp.
11. ***Residual dark drift correction*** – Fits and removes residual "jail-bar"

patterns related to a channel-dependent dark current drift.

12. ***Slope estimation*** – Fits the signal ramp and measures the slope, then estimates the signal in the final, upper plane of the data cube. The output is thus a collapsed, single (upper) plane image, in units of electrons/second/pixel.

13. ***Residual droop correction*** – Removes residual droop by referencing unilluminated portions of the array.

14. ***Short-low stray light correction*** – Corrects for light falling on the SL spectral orders from the peak-up sub-arrays.

15. ***Order cross-talk correction*** – Corrects for optical cross-talk between orders for the SH and LH arrays.

16. ***Flat fielding*** – Applies pixel-to-pixel response correction to the science data.

List of Extended Pipeline Tasks:

- ***Geometric Distortion correction*** – Corrects for the curvature of the spectra in the cross-dispersion direction.
- ***Spectral extraction*** – Extracts spectra from each echelle and spectral order.
- ***Wavelength solution*** – Applies the wavelength solution to each extracted spectrum.
- ***Order tuning*** – Applies polynomial coefficients to correct for diffraction losses and low-order photometric calibration residuals.
- ***Order Splicing*** – Pieces together the orders in the low- and high-resolution science data and corrects for offsets in response between orders and modules.
- ***Flux Calibration*** – Applies a point-source flux calibration to the data, converting electrons/sec to Jy.
- The IRS uses narrow slits and is spatially undersampled (except at the longest wavelengths of each order). In addition, the IRS pipeline is designed to work on point source data. All spectral calibrations are based upon matching point source science data to point source calibration data. Therefore, absolute spectrophotometry with the IRS is very difficult, even in the case of point source science targets. However, the IRS detectors are very stable, showing deviations of less than 1% from campaign to campaign. The overall absolute photometric uncertainty of the IRS is currently 10%, dominated by uncertainty in the models. In addition, there are slope deviations from the beginning to the end of an order with respect to the models of 7% for the high resolution modules and 3% for the low resolution ones. Furthermore, residual features in the flats produce systematic level changes of 2 or 3%. The absolute photometry of the IRS 1-D spectral data is expected to improve with future pipeline releases, but it is recommended that all observers carefully check their spectra against known photometric values and scale the data where appropriate, as they would for all narrow-slit spectra taken with any other ground- or space-based instrument.

7.3.4 Warm/Bad Pixels and NaNs in IRS Data

All four IRS arrays are subject to the damaging effects of solar flares. During IOC, a solar storm deposited in only two days an equivalent dose of protons expected over 2.5 years. This damaged 1% of the SL and SH pixels, and 4% of the LL and LH pixels (mainly the LH pixels). The degradation of the detectors is a continuous process, and it has resulted in a number of "rogue pixels."

A rogue pixel is one with abnormally high dark current and/or photon responsivity (a "hot" pixel) that manifests as pattern noise in an IRS image. The term "rogue" was used originally to distinguish pixels whose hotness was unpredictable, but now rogue pixel masks include those that are permanently as well as temporarily hot. See the SSC website for more details.

As a result of the continual increase in the number of degraded pixels in the IRS Long High module due to the effects of cosmic rays, it is strongly recommended that observations of sources with equivalent point source flux densities of 2.0 Jy or less at 25 μm be accompanied by off-source background measurements. Examination of the properties of the rogues, shows that the rogue dark current can be reduced substantially by subtracting a nearby background spectrum taken within a few hours of the science spectrum. Improvements in the signal-to-noise in the extracted spectrum can be substantial, up to 80% when averaged over an entire order of LH, and when the integration time is split evenly between target and background, since for faint sources the dark current in the rogue pixels is the dominant source of noise.

This recommendation applies chiefly to those observers using IRS Staring mode. Observers who are mapping their sources may only need one additional background observation (if all their map positions fall on the target of interest), or none (if at least one of the map positions falls on blank sky).

To reduce the (increasing) number of rogue pixels, the bias voltage in the Long High module was reduced from 2 to 1.6 V. This happened at the start of IRS 25 (10 October 2005) and resulted in a reduction of the number of rogue pixels by a factor of about 2.5. This comes at the expense of reduced sensitivity in the detector. However, in some cases (very faint sources), the signal-to-noise ratio may be dominated by rogue pixels, and so it will actually increase.

The SSC has developed IRSCLEAN_MASK, available from http://ssc.spitzer.caltech.edu/archanaly/contributed/irsclean/IRSCLEAN_MASK.html, to help in the detection and removal of rogue pixels in IRS data.

Permanently damaged pixels are masked from data processing by assigning "Not a Number" (NaN) status to them. This prevents numerical operations from being affected by their presence. Assigning any other real number would lead to improper processing of the surrounding signal information; for example, during the extraction process. An extraction element containing a damaged (NaN-ed) pixel will obviously contain lower signal than expected, and will appear as an

outlier below the continuum (if a continuum is present). These outliers will be identified by a certain range of values in the 5th column of the extraction (1D spectra) tables.

More information on dealing with bad pixels is available in the Data Handbook.

7.4 Appendix – Definitions

Resolving Power – $R \equiv \lambda / \Delta\lambda$, where λ is the observing wavelength and $\Delta\lambda$ is the full width at half maximum response for the selected entrance slit. Note: at the geometrical optics limit, there are two pixels contained within the projected width of the entrance slit.

Diffraction Limit – $\theta_{dif} \equiv 1.0\lambda / D$ radians when projected onto the sky, where D is the effective diameter of the Spitzer entrance aperture (85 cm). Thus $\theta_{dif} \approx 0.24\lambda$ (λ in µm, θ_{dif} in arcseconds).

Pixel Size – One-half of the width of a diffraction-limited slit at λ_{max}, the long wavelength limit of the associated spectrum. For the SL module, λ_{max} = 14 µm for the purpose of setting the slit width. The exact scaling between pixel size and slit size is module-dependent. The physical pixel size is 75 µm.

f/ratio – The focal ratio (effective focal length/aperture). The f/ratio of the Spitzer telescope is 12.

8 Multiband Imaging Photometer for Spitzer (MIPS)

8.1 MIPS Instrument Description

8.1.1 Overview

8.1.1.1 What is MIPS?

The Multiband Imaging Photometer for Spitzer (MIPS) provides the Spitzer Space Telescope with capabilities for imaging and photometry in broad spectral bands centered nominally at 24, 70, and 160 μm, and for low-resolution spectroscopy between 55 and 95 μm. The instrument contains 3 separate detector arrays each of which resolves the telescope Airy disk with pixels of size $\lambda/2D$ or smaller. All three arrays view the sky simultaneously; multiband imaging at a given point is provided via telescope motions. The 24 μm camera provides roughly a 5′ square field of view (FOV). The 70 μm camera was designed to have a 5′ square FOV, but a cabling problem compromised the outputs of half the array; the remaining side ("side A") provides a FOV that is roughly 2.5′ by 5′. The 160 μm array projects to the equivalent of a 0.5′ by 5′ FOV and fills in a 2′ by 5′ image by multiple exposures. The 70 μm array also has a narrow FOV/higher magnification mode, and is additionally used in a spectroscopic mode. The MIPS cryogenic scan mirror mechanism (CSMM) is integral to all observational operations, allowing selection of different optical trains, image motion compensation during scanned imaging, and one dimensional image dithering. A brief, high-level summary of MIPS for astronomers appears in the ApJS Spitzer Special Issue, specifically the paper by Rieke et al. (2004, ApJS, 154, 25) entitled "The Multiband Imaging Spectrometer for Spitzer." A copy of this paper is available on the SSC website.

There are four basic operating modes and associated Astronomical Observation Templates (AOTs) for MIPS. The Scan Mapping AOT is used to image large areas of the sky in one or more bands nearly simultaneously. The Photometry and Super Resolution AOT is used to image sources smaller than about 2′ diameter, including point sources. The Spectral Energy Distribution (SED) AOT is used to obtain low-resolution ($R\sim20$) spectra in the 70 μm band. The Total Power (TP) AOT is used to measure the absolute brightness of highly extended emission, e.g., zodiacal light. Each AOT includes a number of options, providing good flexibility for designing observations.

MIPS is advancing our knowledge of the far-IR universe by providing the ability to make measurements of unprecedented sensitivity and spatial resolution. Although the design of MIPS was optimized around the four defining science themes for Spitzer, its capabilities were also established from a very general perspective and it is well suited to a broad variety of investigations.

8.1.1.2 Technical Overview

MIPS employs three distinct array detectors (128×128 pixel Si:As BIB; 32×32 pixel Ge:Ga; 2×20 pixel stressed Ge:Ga) to provide high-sensitivity, low-noise, diffraction-limited performance from roughly 20 µm to 180 µm. The short-wavelength Si:As array is combined with a permanent filter to give a bandpass from 20.8 to 26.1 µm with a weighted average wavelength of 23.68 µm. The Ge:Ga array is combined with a filter to give a bandpass from 61 to 80 µm with a weighted average wavelength of 71.42 µm for imaging. This array can also be used in conjunction with a slit and grating to provide low-resolution ($\lambda / \Delta\lambda \approx 20$) spectroscopy from 55–95 µm in Spectral Energy Distribution (SED) mode. The 2×20 stressed Ge:Ga array is combined with a permanent filter to give a bandpass from 140 to 174 µm with a weighted average wavelength of 155.9 µm. The pixel size for all 3 arrays is roughly $\lambda / 2D$ to sample the telescope Airy pattern fully. In addition, a fine pixel scale/narrow FOV can be selected for use with the 70 µm band to provide higher spatial resolution than is possible with the nominal pixel scale, but at a somewhat degraded sensitivity. Selection between the various bands and modes of operation is accomplished through the cryogenic scan mirror, which deflects incoming light into the desired optical path. The scan mirror is also used for image motion compensation during mapping operations, and provides chopping between source and sky to improve the accuracy of the MIPS long wavelength photometry.

8.1.2 Optics

8.1.2.1 MIPS Spatial and Spectral Resolution, and Fields of View

Figure 8.1 schematically illustrates the physical layout of the major optical elements (24 µm Si:As, 70 µm Ge:Ga, and 160 µm stressed Ge:Ga focal plane arrays (FPAs), movable scan mirror, and fixed mirrors and grating) within the MIPS cold assembly. Spitzer's central axis and the telescope focal plane are to the right in this view. Two plane mirrors in the telescope focal plane deflect light into the instrument where it is reflected back by 2 mirrors to form pupils at the two facets of the Cryogenic Scan Mirror Mechanism (CSMM). The CSMM provides chopping at about 0.1 Hz, linear ramp motions, and also deflects the light into the desired optical train. Light is simultaneously sent into the 3 wide-field optical trains, or into the 70 µm narrow FOV train, or into the 55–95 µm SED optical train.

The telescope point spread function (PSF) size ($\approx \lambda / D$) is 6″, 18″, and 40″ full width at half maximum, at 24, 70 and 160 µm respectively. To achieve adequate sampling of the PSF for accurate photometry and for super-resolution imagery (to be produced via post-processing), a pixel size somewhat finer than the conventional Nyquist limit of $\lambda / 2D$ is required. This criterion is met by the pixels of the 24 µm and 160 µm arrays, which have angular sizes of 2.55″ and 16″×18″ respectively. The pixels of the 70 µm array are somewhat larger than

optimal in the wide-FOV mode, having an angular size of 9.98″. This pixel scale was chosen to provide the highest possible 70 μm sensitivity given the expected rate of cosmic ray hits, at the cost of some spatial resolution. In the narrow-FOV mode the 70 μm array pixels have an angular size of 5.2″, small enough to provide excellent performance for photometry and super-resolution, but at the cost of losing some sensitivity.
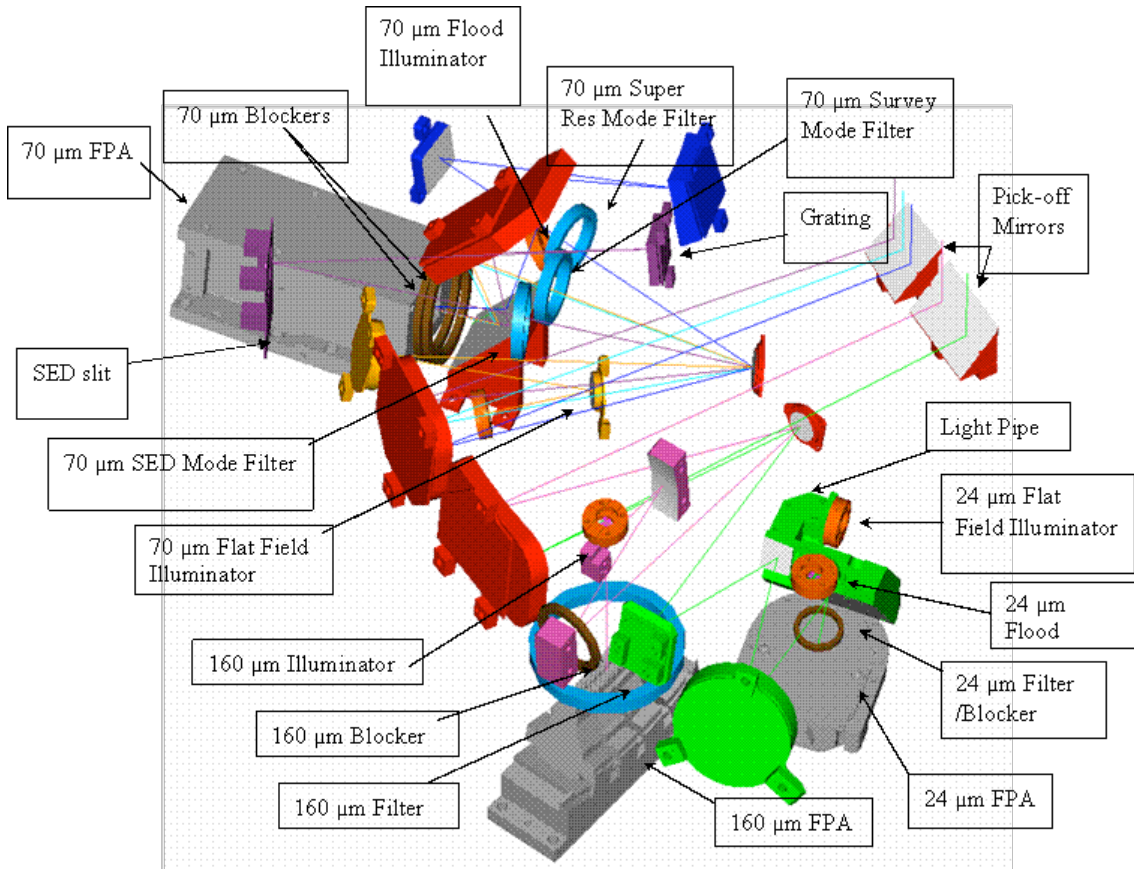
Figure 8.1: Schematic diagram of the MIPS optical train, detectors (Focal Plane Arrays - FPAs), and optical paths. Two mirror facets are attached to the Cryogenic Scan Mirror Mechanism (CSMM): one mirror feeds the 70 μm optical train (normal FOV, narrow FOV and spectrometer/SED), while the second mirror feeds both the 24 μm and 160 μm optical trains. The CSMM provides chopping and one-dimensional dithering for all 3 arrays as well as selecting among the various bands and modes.

All three arrays simultaneously view non-overlapping fields on the sky. The 24 μm array has a ~5′ square field of view; roughly 8′ away is the 70 μm array. The 70 μm array suffers from thermally-activated high resistance in a cable connection that is external to the instrument, but feeds it. This has resulted in significantly increased noise on side B of the array, rendering that side essentially unusable. Also, one readout on side A (located in the lower right corner in data oriented as observers receive it, next to side B) is inoperative; the resulting shape (roughly 2.5′×5′) can be seen in Spot visualizations.

In a single pointing, the 160 μm array views two 5.3′ by 16″ strips separated by a 16″ wide blank strip giving a 5.3′×0.75′ unfilled FOV. The central blank strip of the 160 μm array is filled by taking images at interleaved scan positions. One block of 5 contiguous pixels in the 160 μm array exhibits anomalous behavior attributable to a thermally-activated short in cabling external to MIPS but well inside the CTA. In data oriented as observers receive it, this block is located in the bottom row; it is the third group of 5 pixels in from the left. The "bite" taken out of the array is visible in Spot. These pixels are not quite square at 16″×18″. Occasionally for brevity, we refer to these pixels as 16″ square with the knowledge that they are slightly rectangular.

Table 8.1: MIPS principal optical parameters

Band	Mode	Array Format	Pointed FOV (arcmin)	Pixel Size (arcsec)	λ (μm)	F/#	$\Delta\lambda$
24μm	Imaging	128×128	5.4×5.4	2.49×2.60	23.7	7.4	4.7μm
70μm	Wide FOV	32×32	5.2×2.6	9.85×10.06	71	18.7	19 μm
70μm	Narrow FOV/ Super Res	32×32	2.7×1.35	5.24×5.33	71	37.4	19 μm
70μm	SED	32×24[14]	2.7×0.34[15]	9.8	55–95[16]	18.7	1.7 μm/px
160μm	Imaging	2×20	5.3×2.1[17]	15.96×18.04	156	46	35 μm

The Spectral Energy Distribution (SED) mode utilizes a 3.8'x0.32' slit to produce low-resolution (R=15–25) spectroscopy covering a wavelength range of 52–97 μm. As half of the original array has significantly higher noise, the usable length of the spectroscopic slit is actually about 2'. The SED slit was originally designed to be offset with respect to the array such that an 8×32 portion of the array is unilluminated in SED mode. This strip was supposed to provide a dark measurement for the 70 μm array when taking measurements in TP mode; the other two arrays can be put in the dark by suitable positioning of the scan mirror. This special treatment for the 70 μm array in TP mode is required because of the multiple 70 μm optical paths which might allow light to contaminate the dark measurement. Unfortunately, this region falls on side B of the 70 μm array, i.e., the noisier side. The present TP mode uses side A only, chopping between the

---

[14] Note that full spectral coverage is limited to only 32×12 because of the dead 70 μm readout and the noisy side B.

[15] Slit width is 2 pixels, slit length is 16 pixels; note that full spectral coverage is only obtained over a 2.0′ length of the slit because of the dead 70 μm readout.

[16] Because of a bad readout at one end of the slit, the spectral coverage for 4 columns of the array is reduced to about 65–106 μm.

[17] Includes scan motion required to sample PSF fully. Single frame FOV is 5.3×0.5 arcminutes.

sky position and an internal dark position.

Basic parameters of interest for planning imaging observations with MIPS are
summarized in Table 8.1.

8.1.2.2 Image Quality

MIPS delivers diffraction-limited imaging capability at all three wavelengths due
to the matching of the pixel scales to the width of the Spitzer PSF. In the wide-
field of view mode at 70 μm, the pixels are slightly too large to achieve $\lambda/2D$
sampling of the PSF, so observers who are particularly interested in getting the
highest possible spatial resolution at 70 μm should use the narrow field of view
("super resolution") mode. In addition, image quality at the corners of the 24 μm
field of view is slightly degraded but still close to the conventional diffraction
limit. Empirical PSFs and detailed models of the Spitzer PSF, based on the Tiny
Tim PSF model developed for the Hubble Space Telescope, are available from the
SSC website (http://ssc.spitzer.caltech.edu/). Observers are also referred to the
super-resolution portion of section 8.2.1.2 in this manual for more detail on this
topic.

8.1.2.3 MIPS Spectral Response

The spectral response of the MIPS arrays is given in detail in the following
sections. Blocking of UV through near-IR wavelengths at 24 and 70 μm is such
that for any source that does not grossly saturate the detectors, no measurable flux
passes through the blocking filter. Blocking was tested on-orbit by observing the
point-source modulation transfer function of the Spitzer+MIPS optics, and
checking that only terms corresponding to wavelengths within the spectral
bandpasses are present.

During IOC/SV, 160 μm signals from K stars were detected to be about a factor
of five stronger than expected. Review of the instrument design revealed a
weakness in the stray light control that results in a short-wavelength (1–1.6 μm)
light leak in this band; see section 8.1.3.6.6. Consequences are not reflected in
Table 8.2.

The spectral response of MIPS in the Spectral Energy Distribution (SED) mode
was determined pre-flight using test filters and a far-IR blackbody source. A
tracing of the SED mode spectral response is in Figure 8.5; basic performance is
given in Table 8.1 and Table 8.2. The slit for the SED mode is 2 pixels wide and
24 pixels long; however, 8 of those 24 pixels fall on side B (the noisy side) of the
array. Moreover, another 4 pixels have incomplete spectral coverage due to the
bad readout on side A. The resulting slit length giving complete spectral coverage
is thus 12 pixels. The grating disperses the light parallel to the direction of
motion of the scan mirror; CSMM motions are used to switch between object and
sky positions, while spacecraft pointing is used to move the object between
positions in the slit. The spectrum is dispersed by a reflection grating across 32
pixels, and covers the wavelength range from 55 to 95 μm, giving a nominal

dispersion per pixel-pair of 3.4 µm. The resultant spectral resolution, $\lambda/\Delta\lambda$, is 15 at 52 µm, about 24 at 80 µm, and about 25 from 85 µm to the long wavelength cutoff.
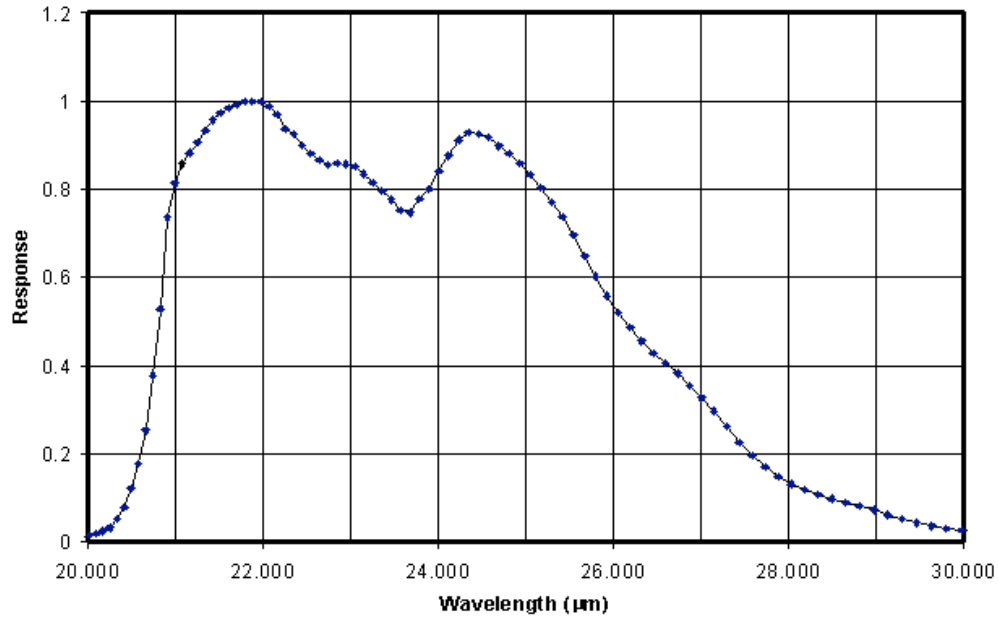
Figure 8.2: Pre-launch response of the 24 µm band including detector and filter spectral response.
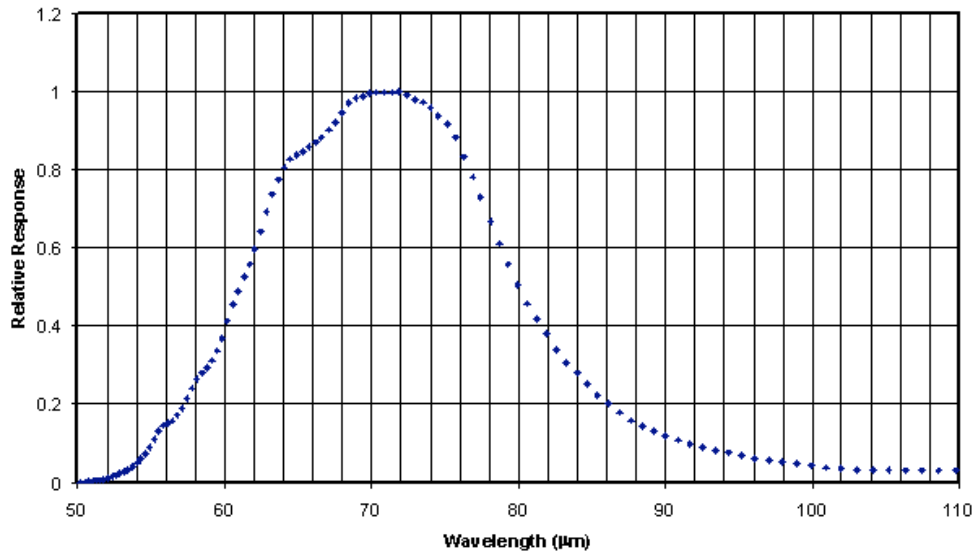
Figure 8.3: Pre-launch response of the 70 µm band including detector and filter spectral response.
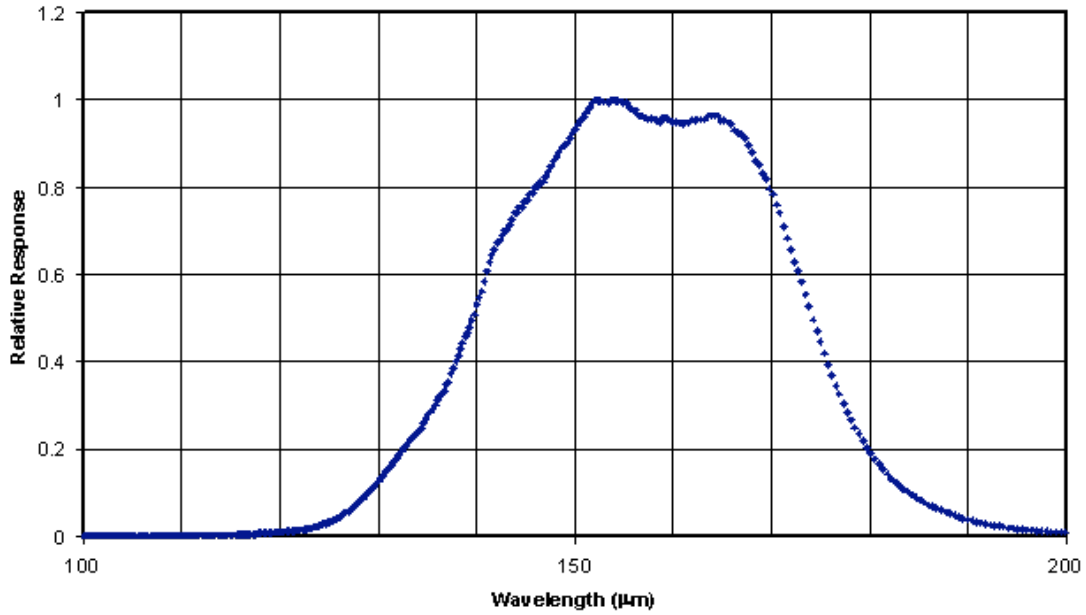
Figure 8.4: Pre-launch response of the 160 μm band including detector and filter spectral response. There is a short-wavelength leak at about 1.2 μm; see section 8.1.2.3.2.
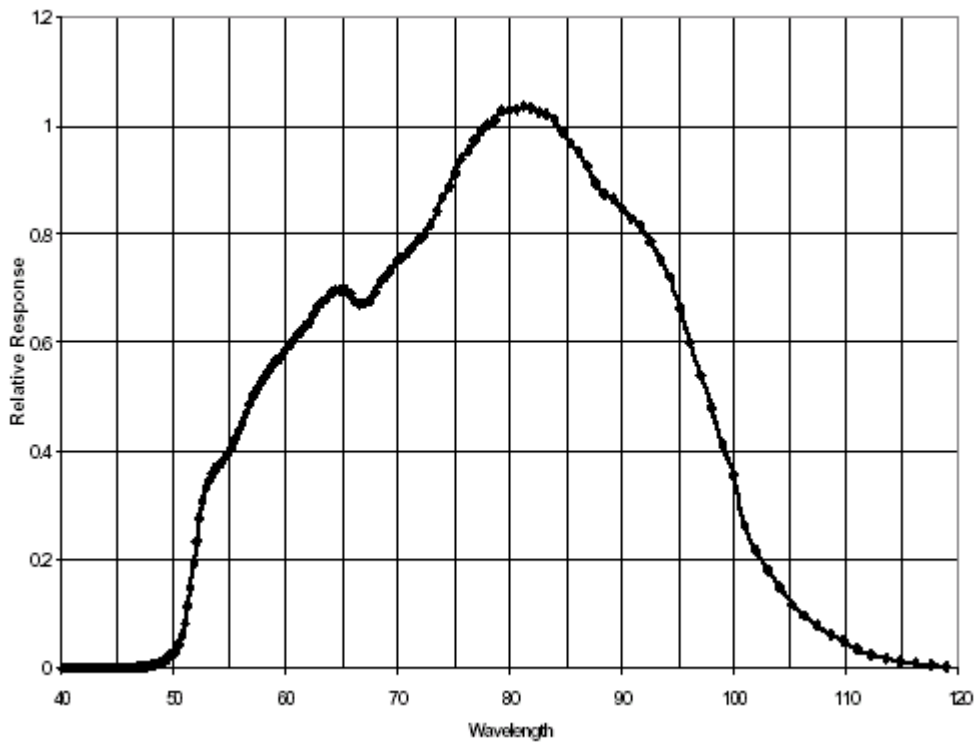
Figure 8.5: Pre-launch SED spectral response including detector, filters, and grating efficiency.

8.1.2.3.1 Spectral Response Tracings

Tracings of the response of the three MIPS imaging bands are included as Figure 8.2, Figure 8.3, and Figure 8.4. Digitized versions of these tracings are available at the SSC website (http://ssc.spitzer.caltech.edu/mips/spectral_response.html). The MIPS Responses, *R*, posted on the web are in units of e- per energy unit.

8.1.2.3.2 Response Parameters

For many applications, it is convenient to characterize the system spectral response in terms of standard parameters. The weighted average wavelengths, λ_{avg}, for the 3 MIPS imaging bands are computed as

$$\lambda_{avg} = \frac{\int \lambda (v/v_0)^{-1} R dv}{\int (v/v_0)^{-1} R dv} \tag{8.1}$$

where *R* is the intrinsic spectral response of the detector, in uits of e- per unit energy (available on the SSC website), v is the frequency, and λ is the wavelength. Basic quantities of the combined detector plus filter spectral response for each of the bands are tabulated in Table 8.2. At 24 and 70 μm, short-wavelength blocking has been verified to be sufficiently effective that signals from a Rayleigh-Jeans spectrum at wavelengths less than one-third the weighted average wavelength of the filter contribute less than 1% of the total.

The response levels in Table 8.2 are given as the product of the filter transmission and spectral responsivity of the arrays, and therefore have units of responsivity, here defined as signal per Jansky falling on the telescope within the projected area of the pixel.

Table 8.2: MIPS spectral response summary

Band	λ_{avg} (μm)	λ_{peak} (μm)	Band-integrated response (e-/sec/mJy)	Cut-on wavelengths		Cut-off wavelengths	
				10%	50%	50%	10%
24 μm	23.68	21.9	730	20.5	20.8	26.1	28.5
70 μm	71.42	71.9	140	55	61	80	92
SED	n/a	n/a	25	53	55	96	100
160 μm*	155.9	152	80	129	140	174	184

* See the discussion in section 8.1.3.6.6 regarding a short-wavelength light leak at 160 μm.

8.1.3 Focal Plane Detector Arrays

8.1.3.1 Si:As (24 μm) Array

The MIPS 24 μm array was developed and furnished to MIPS through the Spitzer InfraRed Spectrograph (IRS) program. It is a Boeing 128×128 pixel blocked impurity band (BIB) Si:As device identical to those in the two short wavelength

spectrometers of IRS except that the MIPS array has an anti-reflection coating optimized for the 24 μm region. The device has four output amplifiers, each of which reads every fourth pixel in a row. That is, the readout is in columns, with an output amplifier for each of columns 1, 2, 3, and 4, repeated for columns 5, 6, 7, and 8. The array is described in more detail by van Cleve et al. (1995; Proc. SPIE Vol. 3354, p. 322-331) and in the IRS Chapter (Chapter 7) of the Observer's Manual (this document), as well as Houck et al. (2004, ApJS, 154, 18). A summary of the characteristics of the flight array is given in Table 8.3.

Table 8.3: MIPS Si:As BIB array performance characteristics

Characteristic	
Detective Quantum Efficiency	~ 60%
Gain (electrons/DN)	5
Departures from Linearity	<10%
Inter-pixel Variability	±2%
Spatial Fill Factor	~98%
Well Depth (electrons)	~3 $\times 10^5$
Dark Current (electrons/s)	3
Read Noise (electrons)	40
Droop Coefficient*	0.33

*See sections 7.1.4 and 8.1.3.3.
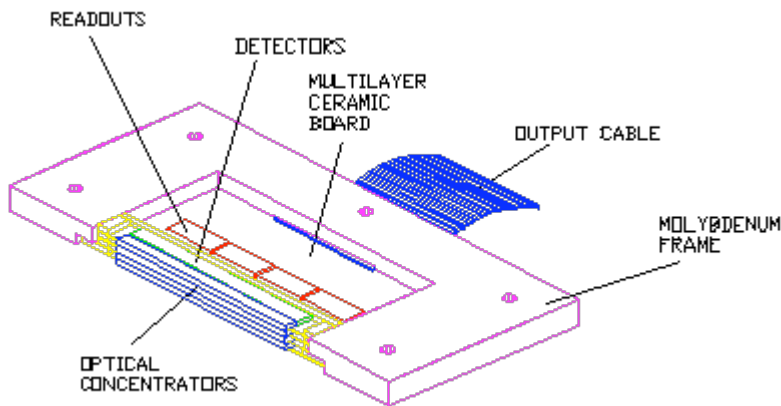
Figure 8.6: 70 μm array 4×32 submodule.

8.1.3.2 Ge:Ga (70 and 160 μm) Arrays

The 70 μm and 160 μm arrays both utilize traditional gallium-doped germanium extrinsic photoconductor detector technology. Because many of the performance characteristics of the two arrays are very similar, the performance of both is discussed below, after the physical layout of the two arrays is described. Differences between the performance of the two arrays are pointed out as needed.

8.1.3.2.1   70 μm Array Design

The 70 μm array is described by Young et al. (1998; Proc. SPIE Vol. 3354, p. 57-

65). Figure 8.6 shows a single 4×32 pixel subarray module from the 70 µm array. Each row of 32 pixels is fabricated from a single 24 mm × 0.5 mm × 2 mm piece of Ge:Ga. The 24 mm length is divided into 32 pixels by photolithography, yielding a dimension of 0.75 mm × 0.5 mm for the illuminated face of each pixel. Absorption of photons occurs over the 2 mm depth of the pixel and is enhanced by a mirror at the back of the pixel that directs transmitted photons back through the sensitive volume. Individual rows are separated by spaces of 0.25 mm. Wedge-shaped germanium optical concentrators are placed in front of each row of pixels to expand the smaller dimension of the pixels to give essentially complete filling of the focal plane and square pixels. Each module contains four readouts, one for each block of 4×8 pixels, mounted on a multilayer ceramic electronics board. A molybdenum frame secures the assembly and the cabling. The vertical dimension of the frame exactly matches that of the 4×32 module, allowing 8 of the modules to be stacked to form the full 32×32 focal plane array (FPA).
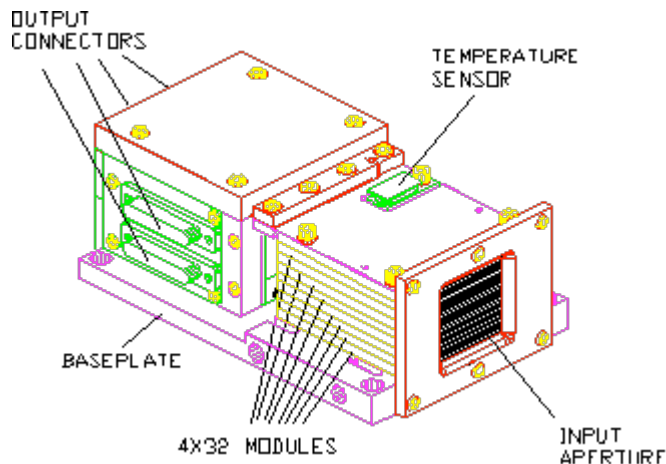
Figure 8.7: Assembled 70 µm focal plane.

Figure 8.7 is a schematic of the full 32×32 70 µm FPA assembly. The 8 modules are stacked and attached to a baseplate. An additional enclosure houses electrical connections and connectors for the camera wiring harness.

8.1.3.2.2   160 µm Array Design

The 160 µm array is described by Schnurr et al. (1998; Proc. SPIE Vol. 3354, p. 322-331). It also has a modular design. Figure 8.8 shows a 5-pixel submodule of the 160 µm array. Each pixel of the array is a discrete block of Ge:Ga 0.8 mm × 0.8 mm × 1 mm in dimension. The 5 pixels are mounted on an alumina electronics board on 3 mm centers, resulting in a very low intrinsic filling factor for this array (see below). Also attached to the board are a single readout for the 5 pixels and an array of 5 integrating cavities.
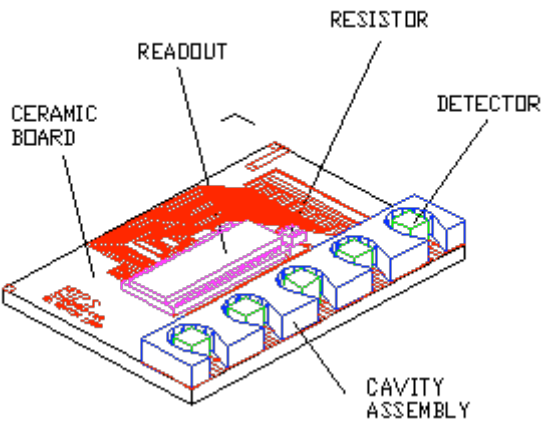
Figure 8.8: 1×5 submodule of the 160 μm array showing discrete pixel elements, integrating cavities, and readout.
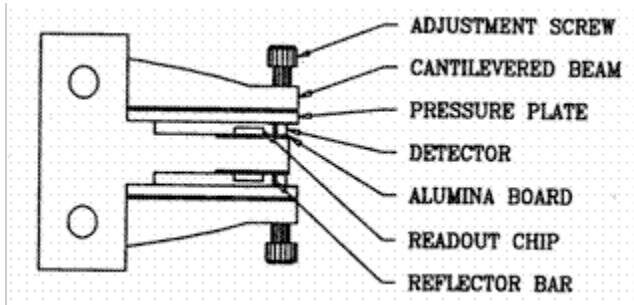
Figure 8.9: Stressing rig for the 160 μm array viewed edge-on.

The pixel material normally exhibits photoconductive behavior only for wavelengths shortward of 120 μm. To extend that behavior to longer wavelengths, the pixels are mechanically stressed to a level of roughly 50% of the material strength. A schematic diagram of the mechanical assembly used to apply and maintain that stress is given in Figure 8.9. The stressing rig contains 2 of the submodules shown in Figure 8.8, oriented back-to-back, and separated by the central metal anvil of the rig. Once the submodules are aligned within the rig, mechanical pressure is applied to the pixels individually, literally by the turn of a screw. A pressure plate between the tip of the adjustment screws and the detector material ensures the even application of force and eliminates the possibility of lateral motion of any of the parts. The amount of force applied is monitored by measuring the conductivity of the pixels as the screws are tightened. Once the appropriate stress level is attained, the screws are mechanically secured to prevent slippage.

Figure 8.10 illustrates the final configuration of the 160 μm FPA. Four of the stress rigs illustrated in Figure 8.9, each containing 10 pixel elements, are mounted to the curved upper surface of a baseplate. Affixed to the top of each stress rig is a 2×5 array of gold plated, light-concentrating, conical apertures.

These concentrating cones expand the effective area of each pixel in the array to ~16″×~18″, and are machined such that each row of 20 has a filling factor approaching 100% in spite of the considerable inter-pixel space required to accommodate the stress rigs and individual integrating cavities. There is a ~16″ wide dead strip down the center of the array which is automatically filled in during observations using multiple exposures and scan mirror motions.
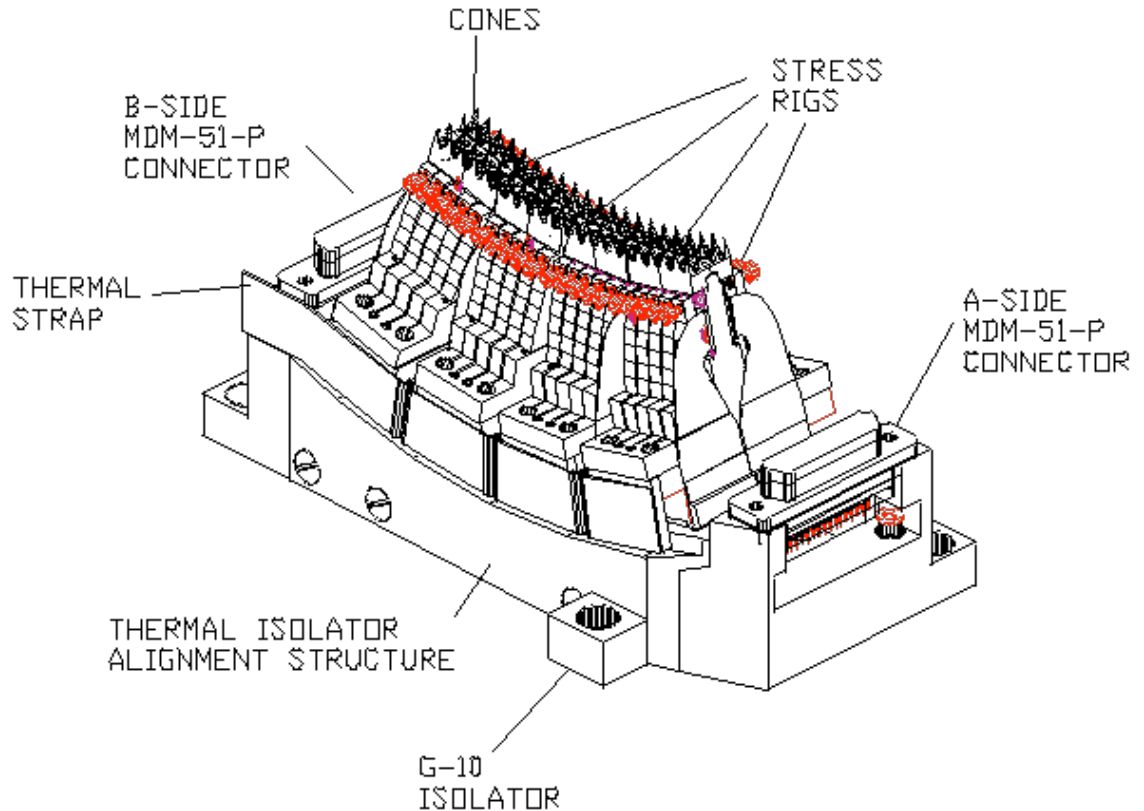
Figure 8.10: Fully-assembled 160 µm focal plane array showing 4 stressing rigs, each containing two 1×5 submodules, and showing the photon concentrator cones.

8.1.3.2.3   *70 µm and 160 µm Detector Performance*

Table 8.4 summarizes the performance of the Ge:Ga arrays as presently known from both lab and on-orbit tests. The response of the 70 µm detectors is linear to about 10%. This high degree of linearity is one of the primary benefits of the capacitive transimpedence amplifier (CTIA) readout design of these arrays, as well as the use of the scan mirror to modulate the signals within the 'fast' response of the detectors (see Detector Behavior, section 8.1.3.3). For the 160 µm array the 'fast' and 'slow' response times are not as widely separated, and the linearity of the response is correspondingly worse than for the 70 µm array. The pipeline corrects for all well-characterized nonlinearities. Saturation for any

particular pixel occurs at $\approx 4\times10^5$ electrons, and is very abrupt due to the use of the CTIA readouts. The saturation value given in Table 8.4 is intentionally smaller than this value to allow for measured pixel-to-pixel variations in responsivity of the arrays, such that an exposure designed using the nominal responsivity values will not saturate any of the pixels in the arrays. Intra-pixel variations in responsivity are uncharacterized, but are expected to be relatively unimportant because the PSF is well sampled. Such variations may be a greater concern for the wide FOV mode of the 70 μm array, where the pixels are relatively large. However, in all cases, the high sampling redundancy in the MIPS data will tend to make the effects of these variations average out. Dead space between pixels is very near zero.

Table 8.4: MIPS Ge:Ga array performance characteristics

Characteristic	70 μm	160 μm
Detective Quantum Efficiency	18%	15%
Responsivity (Amps/Watt)	13.7 Peak 6.9 @ 70 μm	13
Gain (electrons/DN)	7	7
Departures from Linearity	~10%	TBD
Inter-pixel Variability	± 40%	± 40%
Spatial Fill Factor	98%	~ 98%
Well Depth (electrons)*	2×10^5	2×10^5
Dark Current (electrons/s)	156	500
Read Noise (electrons)**	92	280
Radiation Hit Cross-talk	0.5%	<1%

*See discussion in text
** Read noises on orbit are 1000-1500 electrons due to cosmic ray effects.

Cross-talk between pixels that receive a radiation hit during an exposure and neighboring pixels is less than 1%. The rate of cosmic ray hits on the large Ge:Ga pixels is high, about once per 12 seconds. The impact of cosmic ray strikes on data quality is minimized by eliminating cross-talk and by employing a sophisticated scheme for detecting and removing cosmic ray hits in pipeline-processing of the data.

8.1.3.3 Detector Behavior

8.1.3.3.1   Germanium Detectors

All traditional photoconductors, such as the MIPS Ge:Ga unstressed and stressed detectors, show multiple time constant response. In addition, they can exhibit spontaneous spiking and non-monotonic response characteristics (the "hook"). These behaviors are indicated qualitatively in Figure 8.11. Although detailed matching of the theory with observed behavior requires use of numerical techniques, the main features in Figure 8.11 can be understood readily. The initial rapid response is due to the generation and recombination of charge carriers in the

bulk region of the detector. The "hook" response arises through dielectric relaxation effects. The initial rapid migration of free electrons can leave a distribution of charge that reduces the electric field in the bulk and hence reduces the gain of the detector, resulting in a reduction of response after the initial reaction to the signal.
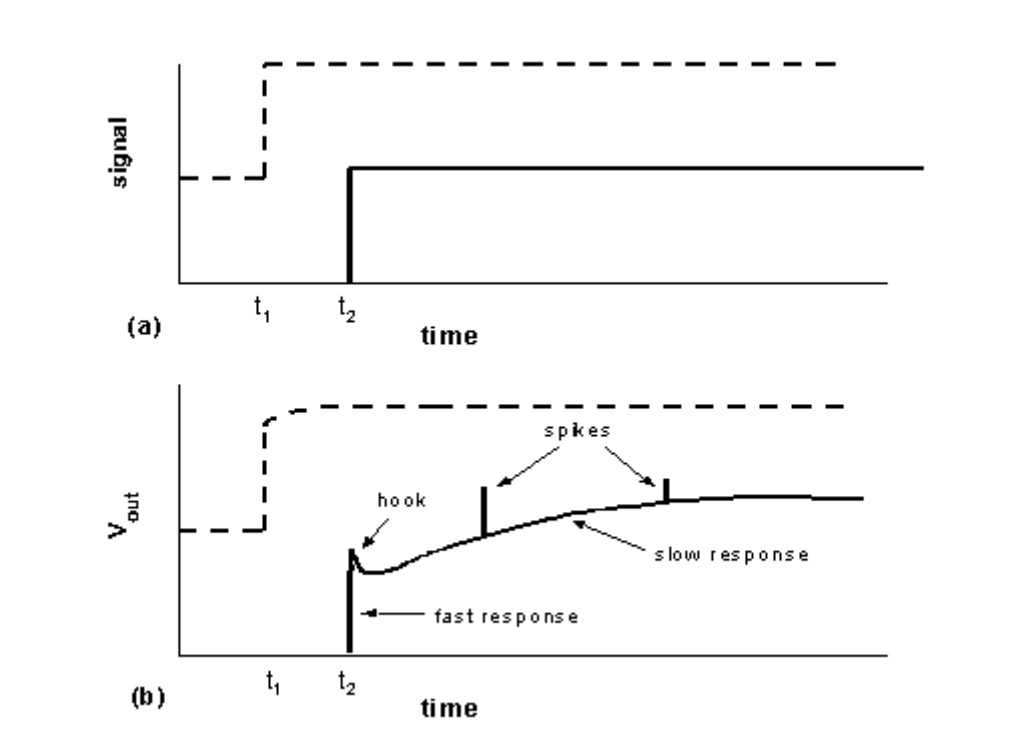
Figure 8.11: Schematic response characteristics of a bulk photoconductor. The signal indicated with a solid line increases from zero to a fixed level at time t_2, as indicated in (a). The detector response, in (b), shows a variety of adjustments including "hook", spontaneous spiking, and a slow tail as a result of the drastic change in detector conductivity. The signal indicated with a dashed line is of the same size as the one at t_2, but the detector was already being illuminated at a level comparable with the signal level, similar to observing a source against a background. The background flux helps suppress the hook and spike response of the detectors, and lessens the impact of the slow response on photometry.

The other effects arise at the detector contacts. In the first order theory of photoconductive response, it is assumed that a charge carrier absorbed at one electrode is "immediately" replaced by one emitted at the other to maintain electrical neutrality in the detector volume. However, the necessity to emit a charge carrier to maintain equilibrium is only communicated across the detector at roughly the dielectric time constant (that is, basically the RC time constant of the detector), so in fact the detector space charge adjusts to a new configuration at this relatively slow rate, not "immediately." The result is a slow increase in detector response lasting tens of minutes (MIPS 160 μm detectors) to hours (MIPS 70 μm detectors). In addition, at the contact junction between conductor and semiconductor, large fields are produced by the migration of charges to reconcile the differing band gap structures. These fields are a result of the

required matching of Fermi levels between the metal contact and the semiconductor detector. To control the peak fields, it is necessary to produce a 'graded contact' by adding dopants through ion implanting to increase the electrical conductivity in a thin layer of semiconductor just below the metal contact. Nonetheless, as the illumination level of the detector is changed, these fields must adjust. During this process, the field local to a small region near a contact can accelerate charge carriers sufficiently to impact-ionize impurities in the material, producing a mini-avalanche of charge carriers that appears as a spike on the output signal. Such spikes are indistinguishable from cosmic ray hits, and are removed in pipeline processing. Care in the manufacture of the contacts for a detector can reduce some of these undesirable effects, such as spiking. It is also helpful to operate the detector at a constant bias; the CTIA readouts used in MIPS have this advantage. Figure 8.12 shows that the response characteristics of the MIPS detectors closely follow the qualitative behavior discussed above.
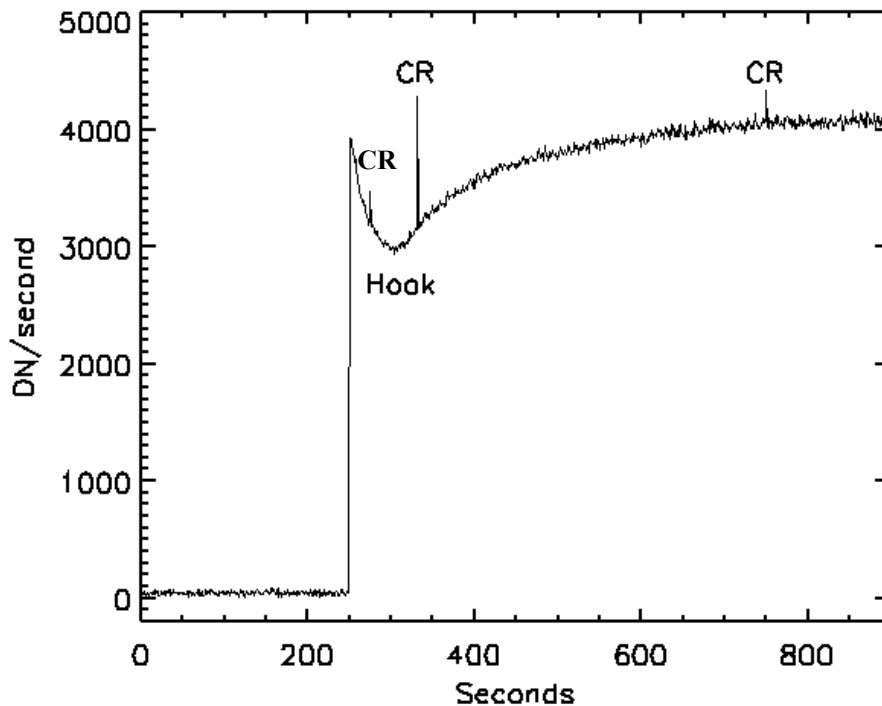
Figure 8.12: Measured response of a pixel of the MIPS 70 μm array to a sudden increase in illumination from a dark background. The hook and long-term increase in response are evident, as are 3 cosmic ray hits.

The fast generation-recombination response (fast response) of the detectors is not subject to the problems inherent in the longer term response described above. MIPS takes advantage of this fact by using the scan mirror to modulate (chop) the source signal on timescales of a few seconds, keeping measurements mostly confined to the fast response regime. At very low background levels such as those provided by Spitzer, it is possible to extract the source signal largely from the fast component, thereby mitigating most of the undesirable longer-term

response effects. Frequent use of on-board calibration sources (see stimulator discussion in section 8.1.3.5) additionally allows tracking of the long-term drifts in response.

One potential area where the long-timescale behavior of the Ge:Ga arrays may be of concern is in situations where there is a high background level in the region of sky being observed. When a significant background flux falls on the detectors, the dielectric time constants are reduced because the background illumination produces a steady state concentration of free charge carriers. At these background levels, maintaining calibration depends on detector response modeling that is being carried out collaboratively by the instrument team at Arizona and the Spitzer Science Center.

8.1.3.3.2   Silicon Array

The 24 μm silicon BIB array is much better behaved than the photoconductors used for the two far-infrared channels. The improvement is inherent in the structure of the detector, which separates the functions of photon absorption and maintaining high detector impedance. Consequently, the infrared-active layer is sufficiently heavily doped that it adjusts to new equilibrium conditions rapidly and avoids the 'two time constant' behavior seen with the germanium detectors. The high absorption in this layer allows it to be small and the cosmic ray hit rate is low. Again, because of the high level of doping, cosmic rays have much less of a tendency to modify the solid state properties of the detector and thereby to produce the problematical response characteristics of the germanium devices.

Nonetheless, the silicon array has some non-ideal behavior. Its output signal is subject to 'droop' whereby the output for a pixel is proportional to the photon signal that fell on that pixel plus a signal proportional to the average signal falling over the entire array. The MIPS flight array has a droop coefficient of 0.33, meaning that 33% of the average integrated charge across the array appears in the output of every pixel. In addition, after observing a bright source, there is a latent image (usually at the < 1% level). The latent images decay very slowly, taking up to ~10 minutes. The droop phenomenon is corrected in the calibration pipeline processing, but **latent images may leave artifacts in the data delivered to users**. Please see the MIPS Data Handbook for additional information and examples.

8.1.3.4 Annealing and Stimulator Use

The response characteristics of the 70 and 160 μm arrays change significantly as cosmic rays cause the slow buildup of residual charge in the detector material. These changes in response are monitored by frequent stimulator flashes (see 8.1.3.5) included in the MIPS observational sequences, at a nominal time interval of 2 minutes. The 24 μm array is of a type that shows minor changes with ionization damage, and calibration with stimulator flashes is a relatively infrequent event for that array.

In addition to tracking changes in the response of the Ge:Ga arrays due to ionizing radiation, the accumulated ionization damage needs to be periodically removed from them. Thermal anneal has proven to be the most effective technique for erasing the effects of cosmic rays. A full thermal anneal is required roughly every three hours, when the heaters raise the germanium detectors to temperatures of ~7 K (70 μm array) and ~5 K (160 μm array). The Si:As focal plane also includes a heater for annealing by raising its temperature to ~ 20 K. However, in this case, the operation is used only when the instrument is turned on to help stabilize the detector.

Unlike stimulator flashes, annealing is a separate operation from normal data taking. For example, anneals may be performed during slews between pointed observations or during downlinks. Observers should expect a significant (but well-characterized) change in the calibration properties of the arrays after an anneal. Although pipeline processing should remove most indications of the change, it is possible that the signal-to-noise ratio on sources will depend on the time elapsed since the last anneal.

8.1.3.5 Calibration with Stimulators

The built-in stimulators (calibration light sources) are the heart of maintaining the germanium detector calibration throughout ground test and on orbit. MIPS has five sets of dually redundant stimulators: 1) flood stimulator for the 70 μm array; 2) flat field stimulator for the 70 μm array; 3) flood stimulator for the 24 μm array; 4) flat field stimulator for the 24 μm array; and 5) flood/flat stimulator for the 160 μm array. The flood stimulators are used exhaustively for calibration purposes. Flat field stimulators offer extra redundancy for the flood stimulators, and are designed to provide a method for monitoring the flat field of the arrays, although operational flat fielding will be based on sky flats. Because the calibration of the germanium detectors is dynamic, the stimulator activity is integrated with the AOTs and repeated flashes are interleaved with the data. The current plan is to use the 24 μm stimulators sparingly, because we expect this array to maintain its calibration well. However, they can be used whenever necessary by issuing a separate engineering command.

8.1.3.5.1   Stimulator Operating Principle

The stimulators operate on the reverse-bolometer principle. That is, they have an emitter that is suspended with a small thermal conductance to the cold sink, and an electrical heater that can heat the emitter above the cold sink temperature.

8.1.3.5.2   Stimulator Usage in MIPS

The flat field stimulator for the 70 μm array is at the focus of a projector that re-images it onto the scan mirror, which is at a pupil in the optical system. The output can be directed through the wide-field 70 μm imaging optical train or through the SED optical train by suitably positioning the scan mirror. The projector is required so flat fields can be obtained for imaging and spectroscopic modes independently (i.e., without light going through both optical trains).

Geometric constraints prevent light from the flat field stimulator from being directed into the high-resolution 70 µm imaging train. The 24 µm flat field stimulator is located in an integrating cavity and its output is conveyed through a machined aluminum light pipe to an exit pinhole at a pupil near the array. The 160 µm stimulator is also in an integrating cavity with a small exit hole in a mirror at a pupil, placed right in front of the concentrators that convey the light to the detectors. The flat fields obtained using the flat field stimulators within MIPS are periodically calibrated against observations of a uniform region of the sky. The flood stimulators at 24 and 70 µm are close to the detectors and broadcast light in a less controlled manner than the flat field stimulators. They are used primarily for relative calibration, while the 160 µm stimulator is used for both calibration and flat fielding.

For observations taken with the 70 and 160 µm arrays, frequent (~2 minute interval) brief (~2 second) stimulator flashes automatically occur during data taking, both during observations of reference flux standards and science observations. The resulting signal from these stimulator flashes provides a continuous measure of the responsivity of each pixel in these arrays. Observers cannot control the rate or placement of the stim flashes. Pipeline processing of the data uses the stimulator flashes to remove long-term changes in responsivity, and provides a relative calibration source for referencing all science observations to flux standards.

8.1.3.6 Calibration Against Celestial Sources

The absolute calibration of MIPS data proceeds in three steps. First, we have determined the best possible calibration of our flux standards at 10.6 µm using other instrumentation. Second, we use "standard" stellar spectral energy distributions normalized to 10.6 µm to extrapolate from that wavelength to the MIPS bands. Additional calibration tasks using asteroids were developed at 160 µm to cope with the ~1.2 µm light leak (see section 8.1.3.6.6). Finally, for the 70 and 160 µm germanium arrays, stimulator flashes provide continuous tracking of responsivity changes since the last observation of a calibration source. Calibration sources are observed at the beginning and end of each MIPS observing campaign to allow the removal of any second-order changes in the calibration of the germanium detectors.

8.1.3.6.1   An Anchor at 10.6 µm

The most accurate absolute calibration near 10 µm is that of Rieke et al. (1985, AJ, 90, 900). This work became the basis for the calibration of IRAS at 12 µm, so the IRAS catalogs propagate it to many stars. However, the fundamental calibrators are best represented in the original paper. Rieke et al. found that calibrations based on extrapolation of A-star spectra disagreed systematically from those obtained in a number of other ways. If we discard the A-star estimates (as was also done by Rieke et al.) but include the other indirect methods that they also discarded, we arrive at fluxes for the three fundamental calibrators in their system as listed in Table 8.5. A more recent absolute measurement of α Boo is in

agreement with the work of Rieke et al. but has substantially larger errors (Witteborn et al. 1999, AJ, 117, 2552). Although not among the Rieke et al. fundamental calibrators, we also include their estimate for Vega, based on intercomparison with the other calibrators.

Table 8.5: Flux densities (Jy) of primary MIPS calibrators at 10.6 μm; rms error of 2% in final adopted value (see text).

	Rieke et al.	Cohen et al.	Adopted
α Lyr	35.3	34.6	34.95
α Boo	655	636	645.5
α Tau	576	569	572.5
β Gem	110	108.6	109.3

M. Cohen and co-workers have published a series of papers presenting calibrated spectral energy distributions of stars, all based upon extrapolated spectra of Vega and Sirius. They use improved models for the atmospheres of these stars, which remove the discrepancies noted earlier by Rieke et al. While their work provides calibrated composite spectra for the stars, here we use only their monochromatic flux densities at 10.6 μm. In Table 8.5, we compare the Cohen et al. and Rieke et al. flux calibrations for the three Rieke et al. fundamental calibrators and Vega. Cohen et al. estimate an overall uncertainty of 3% in their results (Cohen et al. 1996, AJ, 112, 2274). The Cohen values in Table 8.5 reflect a 0.6% upward revision of their fluxes, as suggested by a recent revision in the calibration of Vega at V (Megessier, 1998, Proc. IAU Symp. 189, p. 153), with an accompanying downward revision of their uncertainty to ~2.5%.

The most important result in Table 8.5 is that the two calibration approaches are in excellent agreement for all three calibrators (and Vega). Since the quoted errors for both groups are nearly identical, it is appropriate to take the straight average of the fluxes from the two approaches, and to assign an rms error of 2% to the result.

8.1.3.6.2   *Extrapolation to the MIPS Bands*

To guard against systematic errors, we use three independent methods to extrapolate the calibration to the MIPS bands: 1) Solar Analog method; 2) extrapolation with A star atmospheric models; and 3) extrapolation with semi-empirical models of red giants. Additional tasks at 160 μm use asteroids; see below.

8.1.3.6.3   *Solar Analogs*

Johnson introduced a method of absolute calibration in which direct comparisons of the Sun and calibrated blackbody sources are transferred to a network of "solar analog" stars of type similar to the Sun. The flux offset between the Sun and these stars is assumed to be equal to their brightness difference at some readily measured wavelength (e.g., V band). This method has been used by Campins et

al. (1985, AJ, 90, 896) and further references can be found there. One unique aspect of the application of this method to MIPS is that it is the only feasible "direct" calibration, in which the data can be referred back to a comparison with a calibrated laboratory source at or near the wavelength of interest, without relying on spectral extrapolation. Calibration using the solar-type stars is accomplished by using the empirical fit to the measured solar spectrum (Engelke 1992, AJ, 104, 1248) to extrapolate solar-type spectra through the MIPS bands. Because of the small differences in stellar temperature, and the minimal line blanketing in the far infrared, the Engelke approximation should be very accurate.

8.1.3.6.4 A-Star Model Atmospheres

A-type stars are popular for extrapolating calibrations because it is believed that their atmospheres are relatively straightforward to model. They are probably the most reliable way to extend the MIPS calibration by comparison with theory. A large number of A-stars can be easily measured at 70 μm, and many more (excepting Sirius, which will saturate) can be measured at 24 μm.

The foundation for this approach is accurate atmospheric models, such as those available for Sirius and Vega. Vega itself cannot be used as a calibrator because of its bright debris disk, and a comparison among the measurements of a number of stars will be necessary to identify those without any contamination by disk excesses. Fortunately, enough A-type and solar type stars are available to allow us to reject sources with disk-induced far-IR excesses by simple comparison with the median fluxes of the ensemble of sources.

8.1.3.6.5 Red Giants

Red giants generally have very complex atmospheres and are not favored for extrapolation of calibrations. However, in the far-infrared, their spectra appear to exhibit simple Rayleigh-Jeans-like behavior, as shown by the extensive set of composite spectra assembled by Cohen and co-workers. Thus, the brightest red giants play a critical role in the calibration in the far infrared. Unfortunately these sources saturate the 24 μm array. To cross-calibrate the bright giants with the shorter-wavelength A and G star calibration, it is necessary to observe a group of fainter red giants at 24 μm, and then assume that their spectra are identical to the bright giant spectra for stars of identical spectral type.

8.1.3.6.6 160 μm Light Leak and Calibrations

During IOC/SV, 160 μm signals from K stars were detected to be about a factor of five stronger than expected. Review of the instrument design revealed a weakness in the stray light control resulting in a short-wavelength (1–1.6 μm) light leak in this band. As a result, there was a reworking of the calibration plans for this waveband. Asteroids are now used as the primary celestial 160 μm calibrators.

Asteroids are observed at both 70 and 160 μm (with some also observed at 24 μm), and so the 70 μm channel provides a link to 160 μm via asteroid models.

This is similar to the IRAS 100 μm calibration strategy. Secondary calibrators at 160 μm include leak-subtracted stars, debris disks, extended targets such as galaxies and planetary nebulae, and ISO surveys.

For many types of observations, this light leak does not impact the data at all. Stars fainter than $m_J \sim 5.5$ mag will not be detectable in the leak above the confusion level. No compact extragalactic sources have m_J brighter than 5.5 mag. The leak signal for a star of $m_J = -0.5$ mag is equivalent to ~ 2.5 Jy at 160 μm. The parameter to consider for any given target is the 160 μm /2 μm flux density ratio. For a Rayleigh-Jeans source, this ratio is 0.0001. Anything with a ratio larger than 0.004 will produce uncorrupted data, so most objects will not be affected. Sources with 160 μm fluxes more than a factor of 40 above that of a Rayleigh-Jeans source from $1 - 1.6$ μm will have leak signals <10% at 160 μm. Galactic programs on star formation, ISM, etc., will likely be impacted.

It has been demonstrated that observations of bright point sources strongly impacted by the leak (stars) can be corrected to a high level using an empirical stellar PSF and careful deconvolution (see, e.g., Stapelfeldt et al. 2004, ApJS, 154, 458). Best results will be obtained with observations for which proper planning and data acquisition was made to allow a good characterization of the leak signature for the object brightness and color. For example, plan to obtain identical observations of "calibrator" stars of similar type and brightness to your target stars, but without long-wavelength excesses.

For stellar observations, the strongest signal apparent in a 160 μm observation is likely to be the leak itself. Tests were performed using HD 163588, a K2III star that is routinely observed at 70 μm and has also been observed at 160 μm. Similar results are obtained for HD 36673 (Arneb), an F0Ib star, and for HD 87901 (Regulus), a B7V star. The leak is 15±7 times as bright as the photosphere of a star. Example images are presented in the MIPS Data Handbook.

8.1.3.6.7   Filter Blocking

Using normal stars as calibrators imposes severe requirements on the understanding of near-infrared filter leaks. Leaks were measured by comparing the beam profiles of observations of asteroids and stars. If there had been filter leaks, the stellar point spread function would have included spatial frequencies appropriate to the near-infrared Airy pattern of the telescope, frequencies which physically cannot be associated with a true far-infrared signal. Because the ratio of near-infrared to far-infrared flux is much lower for asteroids than for stars, near-IR filter leaks are negligible for the asteroids. Therefore, a careful comparison of PSFs can reveal filter leaks and allow us to measure them. Spatial power spectra of the PSFs can clearly separate the two frequency components if they are present. This method is described by Kirby et al. (1994, AJ, 107, 2226).

8.1.3.6.8   Dark Current

As part of instrument calibration, dark current must be periodically measured for

each array. This is done at the beginning and end of each instrument campaign, similar to observations of the flux calibration stars. The scan mirror has positions which do not allow light to fall on the arrays, and these positions are used, in part, for dark current measurements. The dark current is measured immediately after an anneal, before the detector responsivity has had time to change due to cosmic ray hits.

8.1.3.6.9 Summary

The MIPS calibration consists of measurements of a number of candidate calibration stars. From these data, we have made a selection of fundamental calibrator stars. At the beginning and end of a MIPS observing campaign, we observe one or more fundamental calibrator stars, establishing the absolute flux calibration for the internal stimulators. Frequent observations of the brightness of the stimulators are used as relative calibrators, allowing the instrument response to be referred to the signals from the celestial standards while only observing those standards occasionally.

Note that additional information relevant to MIPS calibration appears in section 8.2.2, "Performance of the Instrument," and section 8.3.1, "Data Calibration."

8.1.4 Combined Electronics

8.1.4.1 Hardware

To reduce the cost and mass of the Spitzer spacecraft, much of the electronics required by IRS and MIPS are shared, including a warm electronics chassis, power supplies, the instrument processor (RAD6000) and the control and science data interfaces to the spacecraft computer. The IRS instrument has four Si focal plane arrays (FPAs) and MIPS relies on the IRS electronics to control and read out the MIPS Si:As FPA. The block diagram of the combined electronics shown in Figure 8.13 identifies the MIPS electronics, the IRS electronics, and the Common Electronics.

The MIPS/IRS warm electronics are designed to be standby-redundant: two identical electronics boxes are flown. Both boxes are connected to the FPAs and scan mirror in the cold instrument, but only one box is powered at a time. Each circuit that is connected to the cold instrument is designed so that when the circuit is not powered, it does not load the signal line running from the powered electronics box to the cold instrument. To provide further protection from failures, each Ge:Ga FPA is treated as two independent arrays. Independent clock and DC voltages are provided for each 16×32 half of the 70 μm FPA as well as each 2×10 half of the 160 μm FPA.

A total of 40 identical analog chains are provided in the MIPS warm electronics for the two Ge:Ga FPAs: 32 for the 70 μm FPA and 8 for the 160 μm one. The cables connecting the FPA outputs to the MIPS warm electronics are over three meters long. Since these cables are routed from the 1.5 K instrument out to the

300 K warm electronics box, low thermal conductivity wire is used for these cables. Each cable contains 24 pairs of twisted pair manganin wire protected by an outer shield; each output signal is twisted with a complementary ground reference. The total capacitance of the cable, up to 700 pF, must be driven by the array output amplifiers. An instrumentation amplifier in the warm electronics, which helps reject common mode noise on the incoming signal and ground reference lines, receives the signals. The output of the instrumentation amplifier is summed with a DC offset voltage and then fed into a 1 KHz two pole filter. The bandwidth of this filter was chosen to be low enough to limit the FPA noise bandwidth and yet provide a reasonable readout time for the FPA. The outputs of the 40 analog chains are multiplexed down to a single analog line, which is fed to the analog to digital converter (ADC) as shown in Figure 8.13. A 16-bit ADC is used to digitize the processed pixel data from the germanium FPAs, with a resolution of 7 electrons per ADU. The silicon data chain is described in section 7.1.3. For that array, the gain is 5 electrons per ADU, again into a 16-bit converter.
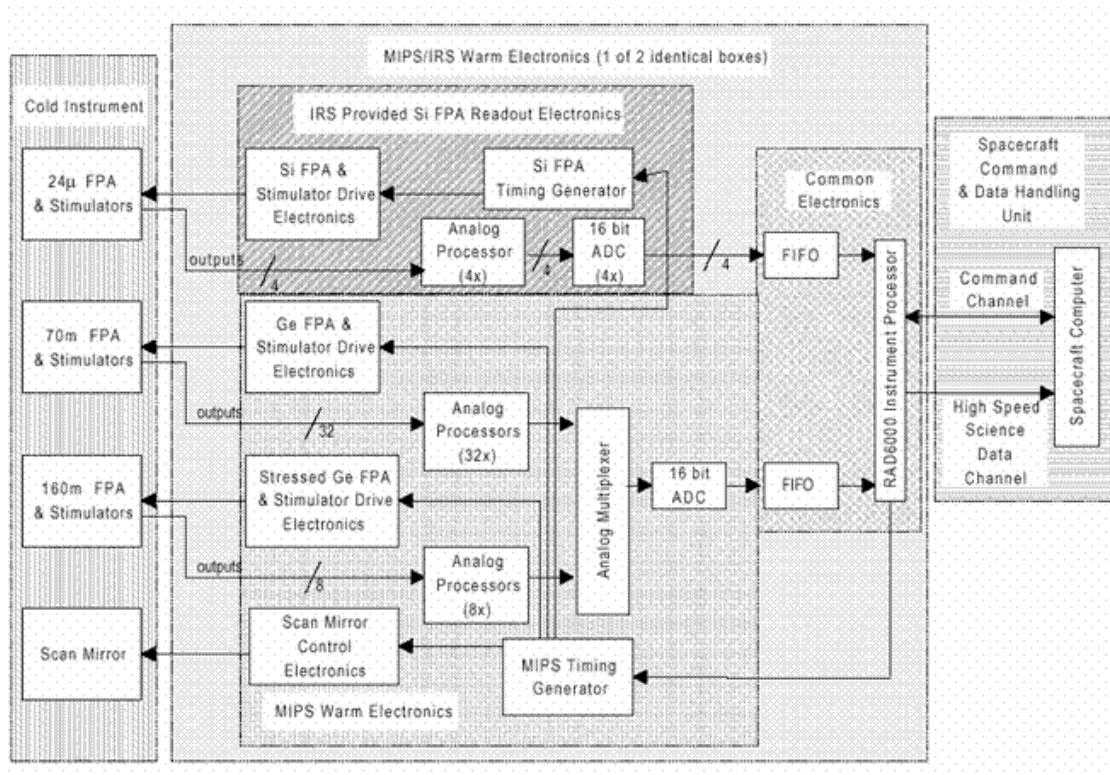
Figure 8.13: Electronic schematic for MIPS/IRS combined electronics package.

As shown in Figure 8.13, a hardware timing generator is used to produce the clock signals for the two MIPS Ge:Ga FPAs, synchronization signals for the MIPS Si:As FPA, and deflection waveforms for the scan mirror. Synchronizing signals are fed from the MIPS timing generator to the IRS timing pattern generator, which produces the actual Si:As FPA clocking patterns. Both of these

timing generators are implemented in hardware with no real time intrusions required by the instrument processor. The MIPS flight software controls the operation of these timing generators through register writes. In operation, the MIPS germanium FPAs are read out at a frequency of 8 Hz and the silicon FPA at 2Hz.

The timing in MIPS has been adjusted to be synchronous with the computer oscillator and other potential sources of synchronous noise. Hence a MIPS second is 5% longer than a conventional one. In this and other sections of this document time information regarding instrument operations (in particular, read intervals and integration times) is **given in 'MIPS time' if not otherwise noted**. In practical terms, the 5% difference is not a concern for observers as it results in slightly longer integration times (and higher sensitivity) than would nominally be achieved. See sections 8.2.2.3 and 8.2.2.4 for specific information regarding the impact on sensitivity and integration time estimates. **Times returned by Spot are in real seconds**, and transparently account for the 5% timing stretch within the MIPS hardware.

Drivers are provided in the MIPS electronics for the FPA calibration stimulator sources and the thermal anneal heaters. The timing generator produces the stimulator flash pulses, which are synchronized with the FPA readout. The thermal anneal heater timing is controlled directly by the instrument processor. The CSMM (scan mirror) is controlled by a type I analog servo system. The actual CSMM deflection angle is continuously monitored and compared with the commanded angle. The error signal between the actual and desired positions is integrated and used to drive the CSMM actuator. Mirror deflection commands are generated digitally by the timing generator and converted to an analog command using a digital-to-analog converter (DAC). The timing generator produces all the required CSMM deflection waveforms, including chop waveforms and the sawtooth pattern required by the MIPS scan map mode. The higher currents required by the CSMM and by the FPA thermal anneal heaters are carried to the cold instrument over a cable constructed of phosphor bronze wire.

8.1.4.2 Software

Figure 8.14 is a top-level block diagram showing the flow of information to and from the MIPS instrument. As shown in the Figure, operational parameters are converted into the appropriate command blocks at the ground station, which are then uplinked to Spitzer. Instrument-specific commands are received and distributed by the spacecraft command and data handler (C&DH). The C&DH provides the interface between the Spitzer spacecraft and the Spitzer instruments. The spacecraft C&DH sends the MIPS-specific commands on to the IRS/MIPS processor; there they are interpreted and executed by the Control Section Flight Software (CSFS). Engineering and science data are gathered by the various components of the MIPS embedded flight software and downlinked, via the spacecraft processor, to the ground for processing.

The CSFS configures and controls the instrument components through four major

electronics sections: the Spitzer Instrument Command and Data handling (C&DH) subsystem interface, the Control Section Electronics (CSE), the Silicon Detector Focal Plane Array (Si FPA) Interface Electronics, the Germanium Detector Focal Plane Array (Ge FPA) Interface Electronics, and the Cryogenic Scan Mirror Mechanism (CSMM) Interface. Within the IRS/MIPS electronics, the CSFS provides the instruments with the computational, interface, and memory resources necessary for performing overall instrument control and science data collection and processing tasks. Commands are accepted through an RS-422 Command/Response Serial interface to the spacecraft C&DH and then processed by the CSFS. The engineering telemetry data are output by the CSFS to the spacecraft C&DH via the same interface. Instrument science data and diagnostic data are output by the CSFS to the spacecraft C&DH via the 1Mb RS-422 Science Data interface. The Si FPA Interface Electronics and the Ge FPA Interface Electronics provide the commanding and science data interface between MIPS detectors and the CSFS. The Si and Ge FPA Interface Electronics also control the instrument stimulators and the Si focal plane array temperature. The Ge FPA Interface Electronics control the MIPS cryogenic scan mirror.
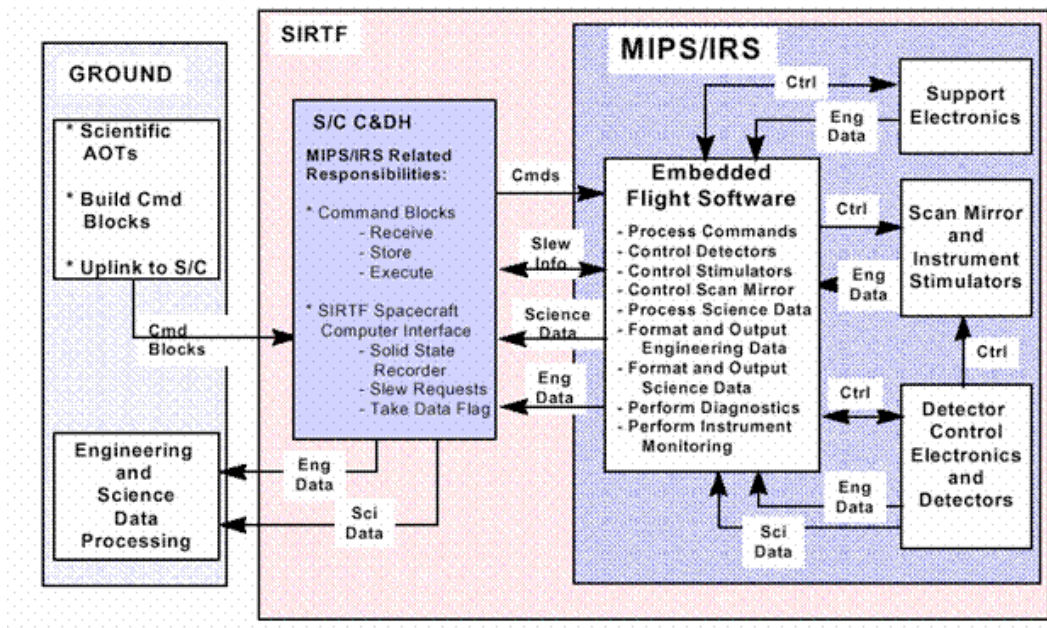
Figure 8.14: MIPS software block diagram.

To perform its main function of collecting and outputting science and engineering data within the constraints of the Spitzer system, the CSFS provides the ability to operate the instrument in two separate software states: the boot state and the operate state. Within each software state, various subsets of the full set of commands and subsets of the full set of telemetry data are valid.

When power is applied to the main electronics of the MIPS instrument, the MIPS processor begins execution of the boot state flight software. The functions for the boot state software are:

- Provide the ability to upload and download data to/from the Error Detection and Correction / Electrically Eraseable Programmable Read Only Memory (EDAC/EEPROM) memory areas from/to the ground system via the spacecraft C&DH.
- Provide the ability to transfer the code from the EEPROM memory area into the EDAC RAM.
- Provide the ability to begin execution of the Operate state software at an uplinked address location.
- Provide the ability to monitor and output a subset of the full complement of engineering data indicating the current state of the instrument.
- Collect detector, memory dump, and diagnostic data telemetry.

To protect this basic upload and transfer capability, the flight software necessary to perform the boot state software functions resides in a write-protected area of radiation-hardened EEPROM. To reduce the complexity and increase the reliability of the boot state flight software, it has been written without incorporating the real-time operating system that is being utilized in the operate state flight software. The boot state flight software utilizes a basic cyclic executive to handle its operations.

After the boot state software has copied the operate state software from EEPROM to EDAC RAM, the instrument may be commanded into the operate state, where instrument science activities can be performed. All instrument calibrations, science observations and instrument diagnostic activities are performed by the operate state software. To accomplish this, the operate state software supports the following tasks:

- Configure and operate the MIPS detectors.
- Configure and operate the MIPS scan mirror.
- Configure and control the MIPS instrument stimulators.
- Collect detector, memory dump, and diagnostic data telemetry.
- Collect and monitor engineering data and scan mirror position data.
- Format and output science data to the RS-422 Serial interface.
- Format and output engineering telemetry data to the RS-422 Serial interface.

8.1.4.3 Array Control

So far as possible within the constraints of data compression and telemetry limitations, each image in the Ge:Ga and stressed Ge:Ga bands is sent to the ground for further processing, which removes cosmic ray hits and other bad data, shifts and adds the frames, and calibrates them to produce the final images of the sky. The Si:As data are fitted onboard by linear regression for each pixel and

only the slopes for each image are sent to the ground (along with a difference of the first two reads). This processing is required to limit the amount of data that must be stored by the spacecraft.

8.1.4.3.1   32×32 Ge:Ga Array (70 µm)

The 32×32 array readouts consist of CTIA amplifiers which are "on" continuously, and whose outputs are sampled sequentially by an on-chip multiplexer. The MUX is cycled at a constant rate to give a read rate of 8 samples per pixel per MIPS second. Each sample is digitized and stored on the spacecraft using lossless compression. Depending on the achieved compression, as Spitzer gets farther from Earth later in the mission, data transmission limitations may require coadding of samples by a factor of 2 at the slow scan rate, and coadding by 2 or 4 samples might be needed at the medium rate. If coadding is found to be necessary, the MIPS electronics will perform that function before sending the data to the spacecraft. The preservation of the detailed data stream for each pixel is critical for this (and the 160 µm) array because of the large hit rate by ionizing particles, as much as 1 hit per pixel per 35 seconds. These hits can only be removed from the data stream in ground processing. The array is reset coincident with a scan mirror flyback, and then reset again half a second later to commence gathering data.

8.1.4.3.2   2×20 Stressed Ge:Ga Array (160 µm)

Operations are similar to the 32×32 array. One distinct difference between operations of the two arrays is that when the observer requests a 10 sec exposure, the 160 µm exposure is split into two 5 sec exposures by resetting the array half-way through the full exposure time. Because of the generally high level of background flux at 160 µm, the array saturates too quickly to be useful in many portions of the sky if a 10 sec exposure is used. An added benefit of performing an extra reset of the 160 µm array is superior rejection of cosmic-ray hits. The extra reset is entirely internal to the MIPS commanding software and instrument, and is transparent to observers. If coadding of samples is implemented (see above), data for both the germanium arrays will be coadded by the same factor.

8.1.4.3.3   128×128 Si:As BIB Array (24 µm)

The 128×128 array also takes data in a mode similar to that just described for the 32×32 array, but some minor differences are required because of differences in its readout, its larger format, and its differing susceptibility to hits by ionizing particles. The array is read out at a constant rate, every 0.5 MIPS sec, and the resulting data sent to the IRS/MIPS computer for processing. However, the data rate is too high to allow sending every frame to the ground. Fortunately, because it has smaller pixels than the 32×32 array (by more than 1000 in surface area), the cosmic ray hit rate is much smaller and it is permissible to process the frames from the multiple reads into a single image (the net integration between resets) before sending it to the ground. The net signal for each pixel in an image is computed by linear regression to determine the current generated by the pixel

between resets. However, there is a loss in dynamic range in this process, since a source that saturates a pixel during the integration will be unrecoverable. To achieve the dynamic range requirements for MIPS, it is necessary also to retrieve the information from the initial 0.5 MIPS sec integration. To do so, the first and second reads of the array are differenced and also sent to the ground.

The readout and resetting of this array proceeds continuously over 0.5 sec, recycling to the same state for the next 0.5 sec. Therefore, the final set of reads for a given integration must start 0.5 sec before the scan mirror flyback. Following reading the array, each pixel is reset during the 0.5 sec read cycle while the scan mirror flies back. A second 0.5 sec reset cycle follows to remove residual signals accumulated during flyback before the scan mirror has settled. Integration on each pixel begins immediately after this second reset. That is, for every scan mirror flyback, there is a one second dead time for each pixel of the Si:As array, and these deadtimes are distributed over a 0.5 sec offset in the order of pixel access and read, with the first pixels beginning integration 0.5 sec after flyback is initiated and the last ones 1 sec after flyback initiation.

8.2 How to Use MIPS

8.2.1 Operating Modes

There are four major operational modes (or Astronomical Observation Templates; AOTs) of the MIPS instrument:

- Scan Map (includes slow, medium and fast scan rates, where incomplete coverage at 160 μm is obtained at the fast scan rate).
- Photometry and Super Resolution (PH/SR; includes super-resolution, large source, and small source options, and can be used to obtain multiple images through the cluster and raster-map options).
- Spectral Energy Distribution (SED; can be used for a single spectrum or step and stare spectral mapping).
- Total Power Mode (TP; used to obtain absolute brightness measurements for highly extended sources).

Observing in any of these modes involves the acquisition of multiple data frames, not just a single frame. The multiplicity of frames provides for rejection of cosmic rays, calibration of the Ge:Ga focal plane array data in the light of the two-time-constant behavior of those detectors, adequate sampling of the point spread function (PSF; especially for super resolution observations), and, in the case of the 160 μm array, for building up a filled image using multiple offset exposures of that 2×20 pixel array. The total number of images obtained depends on the total integration time needed, the observational mode, and the array. The multiplicity of frames is prescribed within the Astronomical Observing Templates (AOTs), and **the observer selects the number of times (cycles) to repeat the basic pattern of a template in order to build up the desired integration time.** The exposure times for each frame are also limited by the AOTs, with the result that the total integration times that can be specified are quantized. The details of

this are described in the following sections and in section 8.2.3 (which also includes a collection of "Best Practices").
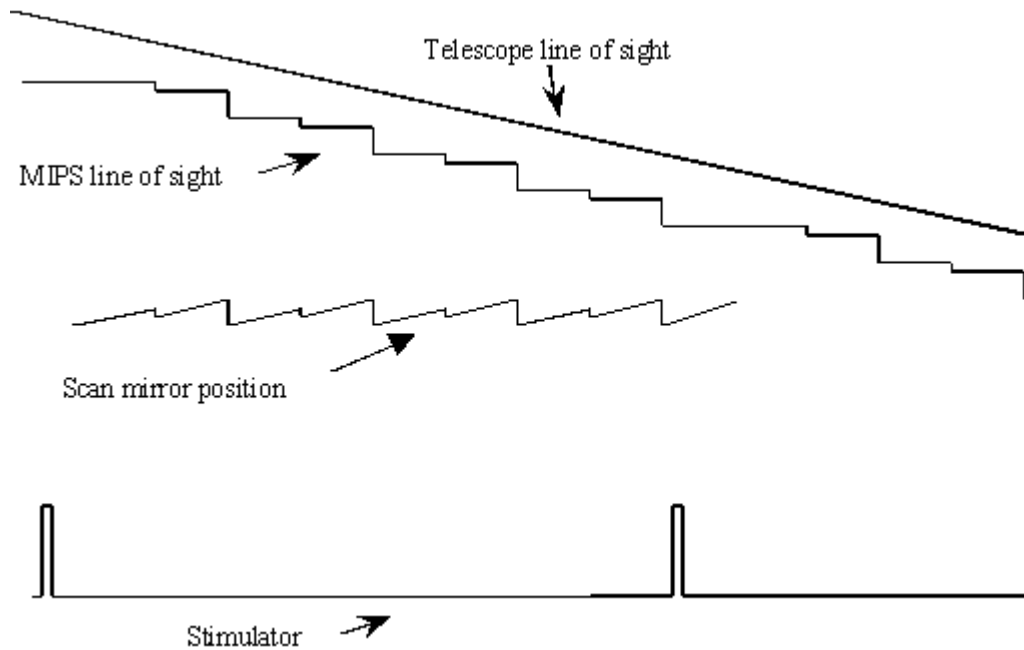
Figure 8.15: Schematic representation of synchronization of scan mirror motions, telescope scan, and stimulator flashes.

8.2.1.1 Survey Scan Map Mode

The scan map mode is designed to provide efficient mapping of large areas on the sky. A ramp motion of the scan mirror compensates for continuous telescope scanning motion, freezing the images on the arrays. The scan map mode avoids having to repoint and stabilize the telescope between exposures. A schematic representation of this is shown in Figure 8.15.

Redundant images greatly improve the reliability of MIPS data, especially from the germanium arrays (e.g., to help remove cosmic ray effects, or uncertainties in the calibration of highly extended emission in a map). A single scan pass (or "leg") provides multiple redundancy at 24 and 70 μm (Figure 8.16). However, single scan legs provide only single redundancy at 160 μm at the slow and medium scan rates, and only ½ *coverage* at the fast scan rate (Figure 8.17 and Table 8.6). Redundant coverage at 160 μm, which is strongly recommended, can only be obtained if the observer specifies at least two overlapping scan legs; the SSC strongly encourages observers to obtain at least 4 images of a 160 μm target if the 160 μm data are important to the science goals. At 24 μm, another scan map at a later time is needed to identify asteroids; most asteroids at 70 μm should be identifiable by the ratio of their 24 and 70 μm colors. At all bands, a second map is desirable to remove transient events.
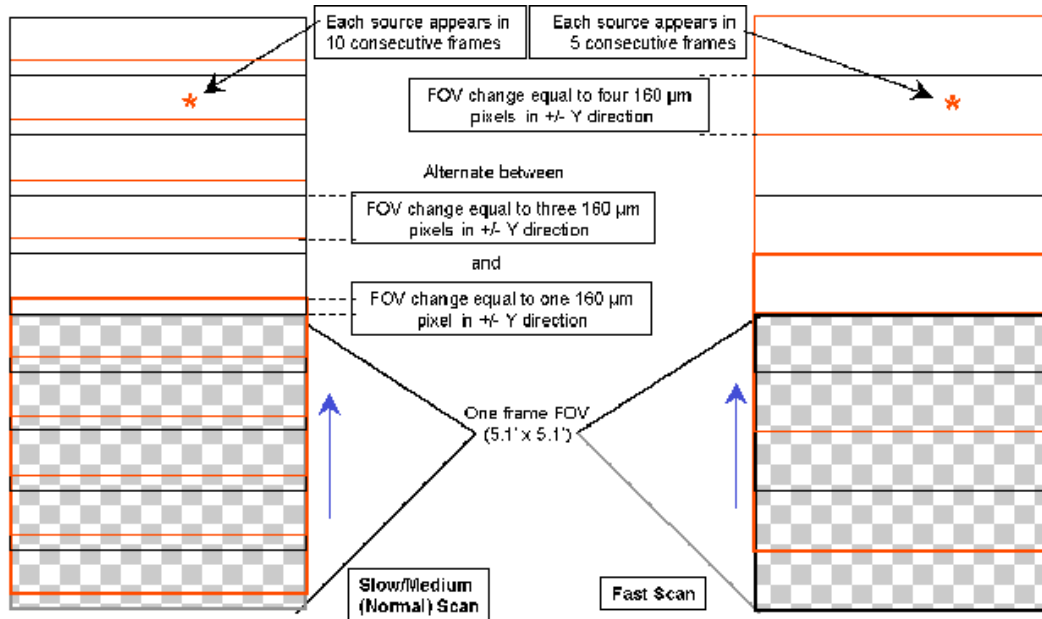
Figure 8.16: Operation of 24 and 70 μm arrays during scan mapping. During normal operation, every source is observed ten times in a single pass; during fast scan, each source is observed five times.

Observers are cautioned that the Scan Map AOT does not intrinsically provide any data redundancy at 160 μm at any scan rate. Those who wish to have high-quality scan data at 160 μm must take care to specify repeat coverage of the map area, either using sub-array scan leg offsets or repeat visits to the map area. See the discussion of data redundancy for the 160 μm Photometry AOT, section 8.2.1.2.

As a result of the on-orbit realities with the Ge detectors, for full sky coverage at comparable signal-to-noise, **observers should step by at most half-array widths (2.5′) in the cross-scan direction between scan legs.** If 160 μm data are required, observers should be sure to design their observations to include sufficient redundancy at all locations in space, i.e., step by no more than 2.5′ and preferably 2′ to avoid the inoperative readout.

8.2.1.1.1 Telescope Motions

The telescope performs a scan at a constant rate. The observer can select from three discrete rates — slow, medium, and fast (approximately 2.6, 6.5, and 17.0 arcseconds per second). A *scan leg* is defined to be a single scan pass in one direction (forward or reverse). A scan map observation always starts in the "forward" direction (in the +Y direction in the Spitzer focal plane; see Figure 4.5). The *scan length* is defined to be the distance along the scan track for which coverage is obtained with all three MIPS arrays. Scan lengths can be as short as 0.5° or as long as 6° (depending on the scan rate). There are regions of overscan at each end that have coverage by some of the arrays but not all. During the full-coverage portion of the observation, the telescope position should be within at

least 1.3″ (rms) of the expected position along the scan track. At the end of a scan, the spacecraft pauses, then begins a new scan in the opposite direction to the first, with the same scan rate and either along the same track or offset by anywhere from 5″ to 5′ in the cross-scan direction, as selected by the observer.

For typical scan maps, scan legs can be assumed to be linear. Very long scans depart slightly from linear tracks. This curvature is determined by the offset between the arrays and the telescope centerline (X-axis), and any misalignment between the scan mirror rotation axis and the spacecraft Z-axis (Solar direction; see Figure 4.5 and Figure 3.3). Also, the scanned region on the sky is in general not exactly 90º from the line to the Sun, but can be anywhere within a 40º wide annulus (see Figure 3.5). The net effect of these terms on the scan leg geometry is that the dominant direction of scan motion is a great circle centered on the spacecraft Z-axis. The motion is made such that the rotation is around a point 90º from the telescope boresight, and within 2º of the great circle centered on Spitzer and going through the center of the Sun. The pointing of the boresight is such that the scan leg centerline is positioned through the array centers. Perturbations that amount to small circle deviations are required to account for the focal plane layout and the actual scan mirror rotation axis relative to the Z-axis. The resultant curvature in the scan leg is <1% over a 6º scan. The relative pointing accuracy between two consecutive scan legs is expected to be ultimately better than 3″ (radial) for at least 4000 seconds time.

During the course of an AOR, the Spitzer focal plane orientation on the sky is fixed. Therefore, there is no rotation of individual scan paths, nor is there any relative rotation between scan legs obtained within a single AOR. This holds in spite of the orbital motion of Spitzer during the period of the AOR. Scans obtained in separate AORs likely will have different orientations. The magnitude of the rotation depends on ecliptic latitude and time between AORs. To minimize such variations in scan orientation, observations may have to be constrained such that they are done within a short period of time, particularly when the targets are far away from the ecliptic. (See section 5.5.3 for more about constraints.)

8.2.1.1.2   *Instrument Functions*

At the start of the scan, the scan mirror is positioned near its "center" position, where all three detector arrays view the sky; data are collected from all three arrays. The scan mirror is moved in a one sided ramp with "flyback" around this position; during the ramp motion, the scan motion of the telescope is compensated by the motion of the scan mirror to freeze the field on the arrays. This operating mode is also called "freeze frame scanning." The flyback occurs at intervals of 1 and 3 pixels of the 160 μm array alternately (i.e., ~16″ and 48″) at the slow and medium scan rates, and every 4 pixels worth of slew at the fast rate (i.e., ~64″). When the telescope scan direction is reversed, the parity of the scan mirror ramp/flyback is reversed. The pixel sizes for the stressed Ge:Ga (160 μm) and Si:As (24 μm) arrays provide good sampling of the telescope PSF. Improved sampling on a fractional pixel level in the scan direction is obtained on the 70 μm

array as the source transits, because the frames are taken at non-integer shifts at its pixel scale. In the cross-scan direction, a slight difference between the slew direction and the scan direction also provides fractional pixel sampling of the PSF.
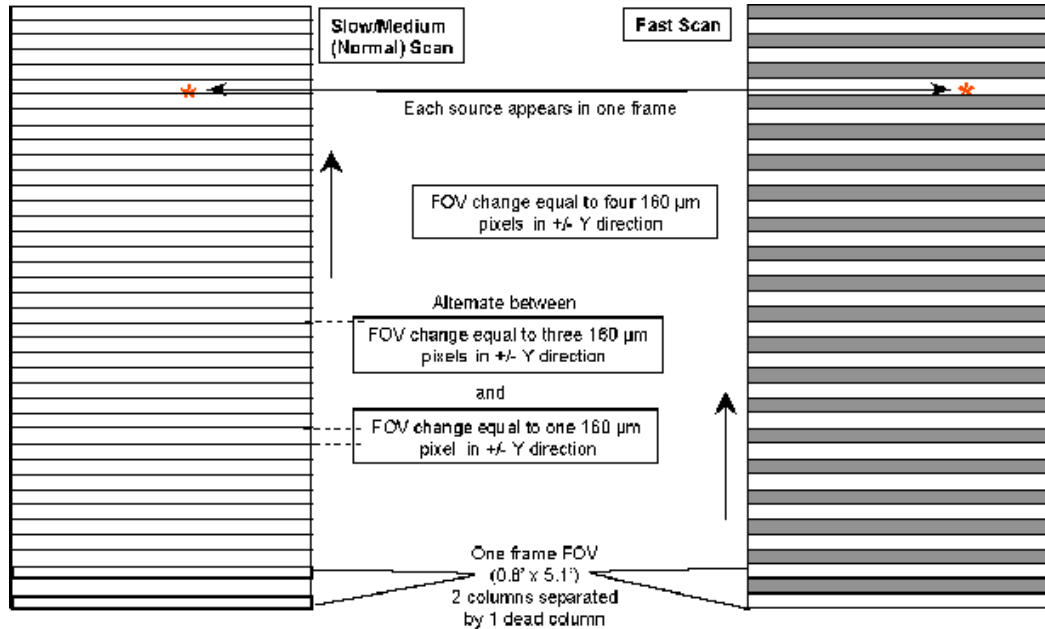
Figure 8.17: MIPS operation at 160 µm in scan map mode. In normal operation, every point along the scan track is observed once; redundant coverage requires rescan. In fast scan mode, the sky coverage is in bars, one pixel wide (16.0″) with every other bar not observed. Full sky coverage requires rescan to place the pixels in the uncovered areas on the rescan.

Frequent stimulator flashes are required to calibrate data from the two germanium focal planes (see section 8.1.3.5). To allow scans that take longer than the time between germanium stimulator flashes (~2 minutes), the motions of the scan mirror and telescope are coordinated. The scan mirror gradually "gets ahead" of the telescope, so that the individual frames lead the telescope by increasing amounts as the telescope scan progresses. This lead grows by linearly offsetting from the zero point of the scan mirror motion. When the time arrives for a stimulator flash, the scan mirror resumes its action at the zero point, and the time for one exposure has been made available for the calibration without loss of sky coverage. This process also guarantees that the array is pointed at a previously imaged portion of the sky during the stim flash, allowing for accurate subtraction of the sky signal. This pattern of operation is illustrated in Figure 8.15. The resulting coverage on the sky is shown in Figure 8.16 and Figure 8.17.

Note that the first several frames of a scan leg necessarily have extrapolated stim backgrounds (see Data Handbook for more information on this effect). In general, but especially in the case of fast scans over bright objects, be sure to have enough background measurements around your source that extrapolated stims are not an issue.

8.2.1.1.3   Fast Scan Mode

Instrument operations at the fast scan rate (nominally about 17″/sec) are slightly different from those at the lower rates. Because the efficiency would be undesirably low at this scan rate if full sampling were obtained at 160 μm, that option is not supported and full coverage at this wavelength requires a second scan. The scan mirror flyback is set to four 160 μm pixel widths (64″), with no alternation of flyback amplitudes. As a result, there is alternating one-pixel-wide gaps in 160 μm maps acquired in fast scan mode. Given that the pixels are 0.4 λ / D wide and separated by the same amount, any reasonably bright sources will be "sighted" at 160 μm even in the partial coverage obtained in a single pass. A general requirement for reliable scan maps is that the same area be scanned a second time (at a later date/time) to screen out asteroids and other transient phenomena. If a rescan is carried out very soon after the first, it can be approximately interleaved such that both scans together provide a nearly full 160 μm map.

Note that the SSC strongly encourages all observers to obtain at least 4 images of a 160 μm target if the 160 μm data are important, e.g., a forward and back scan leg over the same area with 2 map cycles, and with a 5-pixel offset in the cross-scan direction for fast scan (5 px compensates for the inoperable readout; see section 8.1.2.1).

Table 8.6: Total single-pass MIPS scan map integration time per pixel and overscan time

Scan Rate	24 μm (sec) [time (sec) per frame]	70 μm (sec) [time (sec) per frame]	160 μm (sec)	Overscan time* (min)
Slow (2.6″/sec)	100 [10]	100 [10]	10	1.2 min
Medium (6.5″/sec)	40 [4]	40 [4]	4	3.2 min
Fast (17″/sec)	15 [3]	15 [3]	3 **	7.9 min

Note: Each source appears in 10 consecutive frames for slow and medium scans at 24 and 70 μm, each source appears in 5 consecutive frames for fast scans at 24 and 70 μm, and each source appears in one frame at slow and medium scans at 160 μm. See Figure 8.16 and Figure 8.17.
* See section 8.2.2.9.2.
** Only ½ of map region is covered by the 160 μm pixels at the fast scan rate in one scan leg.

The total integration time per point on the sky per pixel in a single scan leg is provided in Table 8.6; overscan times are discussed further below in section 8.2.2.9.

8.2.1.2 Photometry Mode

Although both photometry and super resolution are supported in a single AOT, for illustration we describe the procedures separately. The photometry mode is designed to obtain multi-band observations of reasonably compact sources; sources that are too extended for the procedures described below should be imaged using MIPS scan mapping instead. Super resolution emphasizes thorough sampling of the point spread function through a combination of small pixels relative to the Airy disk and of sub-stepping of the image relative to the pixels. Special observing modes are used to achieve this goal at 70 and 160 μm, and at 70 μm a separate optical train providing a finer pixel scale is also employed. **Note that super-resolution data requires additional post-processing (beyond what the SSC currently provides) to extract higher spatial resolution.**

In the following descriptions, we give overly specific examples, for example with particular pixel addresses for where the image falls. We do so for purposes of illustration only. Actual placements on the arrays may not be exactly those given here, but relative placement remains. Observers are referred to Table 8.11 below for a summary of photometry mode integration times. Note also that the multiplicity of frames is prescribed within the AOTs, and the **observer selects the number of times (cycles) to repeat the basic pattern of a template in order to build up the desired integration time.**

There are multiple options within the photometry AOT. "Compact source" can also be referred to as "small field," and "large source" as "large field." Additionally, for 70 μm, "fine scale" pixels can also occasionally be referred to as "narrow field," and "default scale" can be described as "coarse scale."

8.2.1.2.1 Photometry – 24 μm, Compact Source (Small Field)

One cycle of the basic 24 μm photometry observation acquires 14 separate images of a source. The locations of these 14 images are illustrated in Figure 8.18 and listed in Table 8.7. To start the observation sequence, the source is centered to the left of the array center by 25 pixels (65″) and the scanning mirror is chopped in a nearly symmetrical fashion, above and below the array center, to obtain a column of 7 images. The spacecraft is then slewed to position the source 25.5 pixels (66.3″) to the right of center and another column of 7 images is obtained. Table 8.7 contains the resulting approximate image positions in units of pixels, with the origin taken to be the center of the array. Array distortion and the twist of the scanning mirror with respect to the array columns have not been taken into account in this table, but the positions are good enough for observation planning purposes.

The observer specifies only one parameter, the number of observing cycles. If the observer specifies N cycles, the sequence repeats the first column N times before the spacecraft moves to take N sets of data in the second column. There are two complications to this simple scenario: (1) two extra exposures are obtained at positions 1 and 8 for each AOR, and (2) exposures 1 and 8 of the first cycle (and a

few subsequent cycles in long AORs) are one second shorter than the specified exposure time. Spot takes this all into account when it calculates the total integration time, but it is approximately $((N \times 14)+2) \times$ (Exposure Time)-2 seconds. The 2 sec reduction comes about because the first second of frames 1 and 8 are used for obtaining calibration data, not science data.
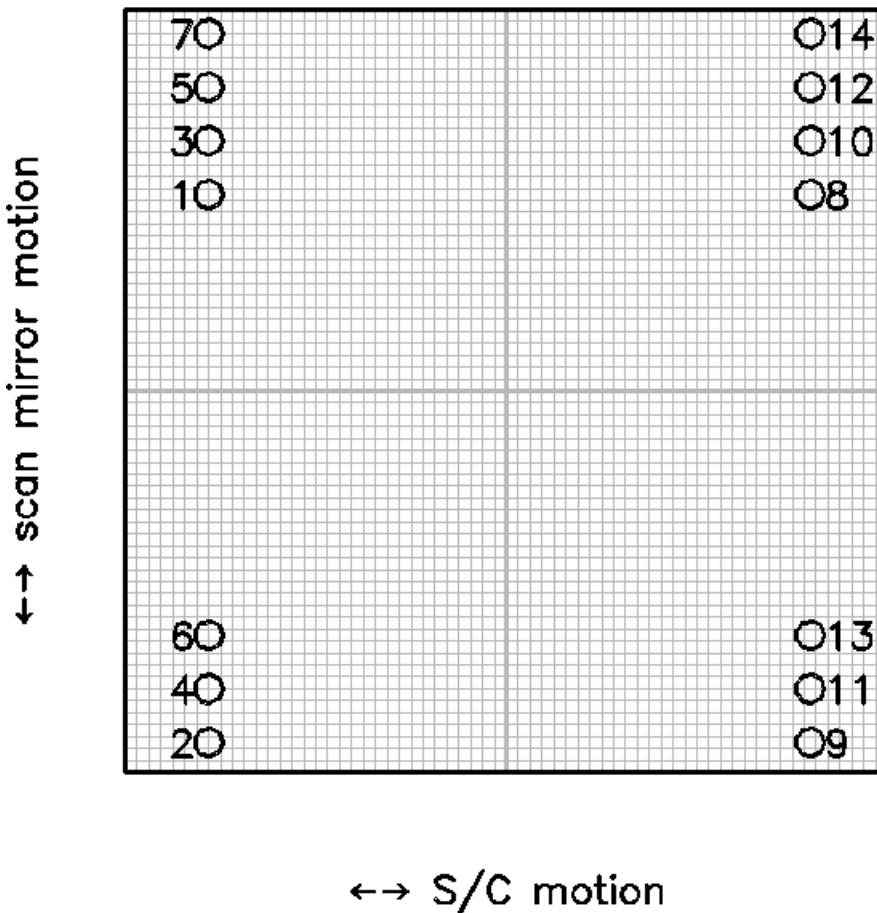
Figure 8.18: Photometry/Super Resolution for compact sources with the 24 μm array. The positions of images in a single cycle of a 24 μm photometry observation are shown relative to the central 64×64 pixels (shown schematically by the grid) of the full 128×128 array. The circle diameters are the FWHM of the Airy disk. The first set of 7 images are dithered using the scan mirror, then the spacecraft is offset by a half-integer number of pixels, and the second 7 images are acquired. Frames 1 and 8 of *the first cycle* of an AOR (and a few subsequent cycles in long AORs) are 1 sec shorter than the observer-requested exposure time; the other 12 frames of such cycles are taken with the full observer-specified exposure time. Not shown are 2 additional frames, taken at the positions 1 and 8, obtained for each AOR.

Note that half of the images are obtained at each spacecraft pointing. The telescope relative offset accuracy allows meaningful fractional pixel offsets. Moreover, in practice, alignment tolerances make the chop motion slightly non-

parallel to the array columns. In any case, the implemented pattern provides a well-sampled Airy disk. Although for simple photometry, the images can be combined by integer shifts, careful processing will be required to take full advantage of the oversampling implicit in these images.

Table 8.7: Source positions for 24 μm compact source photometry (in units of pixels, the origin being the center of the array).

Frame number	X position	Y position
1	-25.0	16.7
2	-25.0	-30.5
3	-25.0	21.3
4	-25.0	-25.9
5	-25.0	25.9
6	-25.0	-21.3
7	-25.0	30.5
8	25.5	16.7
9	25.5	-30.5
10	25.5	21.3
11	25.5	-25.9
12	25.5	25.9
13	25.5	-21.3
14	25.5	30.5

Note also that the final highest-S/N region of the combined mosaic is not a full FOV, but rather a ~3′ square area.

The photometry sequence is currently implemented as described here, but the number of images obtained per cycle may change without notice later in the mission. The basic technique will always be to define an image pair to be produced by a scan mirror chop of slightly less than half the array width, and to offset successive pairs by a small amount to improve the pixel sampling.

8.2.1.2.2 Photometry – 24 µm Large Source (Large Field)

For larger ($\geq 2'$) sources, dithering purely on the 24 μm array is no longer desirable. Instead, the scan mirror is used to obtain a series of images separated by only a minimal amount and on half-pixel centers; see Figure 8.19 and Table 8.9. The telescope is then redirected to a position $> 5'$ away from the source, and the sequence of frames is repeated to obtain a sky image. The offset to this "sky" position is in the scan direction, and the observer specifies the magnitude of the offset. The nominal pattern for these observations is a 1×5 set of images in-scan, separated by 1.5 pixels = 3.825″ for each, followed by a cross-scan spacecraft offset of 4.5 pixels (11.475″) and a repeat 1×5 set of images. In this sequence, frames 1 and 6 are 1 sec shorter than the observer-specified exposure time at both

the object and sky positions, and two extra exposures are obtained at object and sky positions per AOR, similar to what happens in the compact source photometry observation sequence described above. As above, the observer specifies only one parameter, the number of observing cycles. The total integration time per pixel in the final images is very nearly equal to $((N\times10)+2)\times$(Exposure Time) -2 seconds.
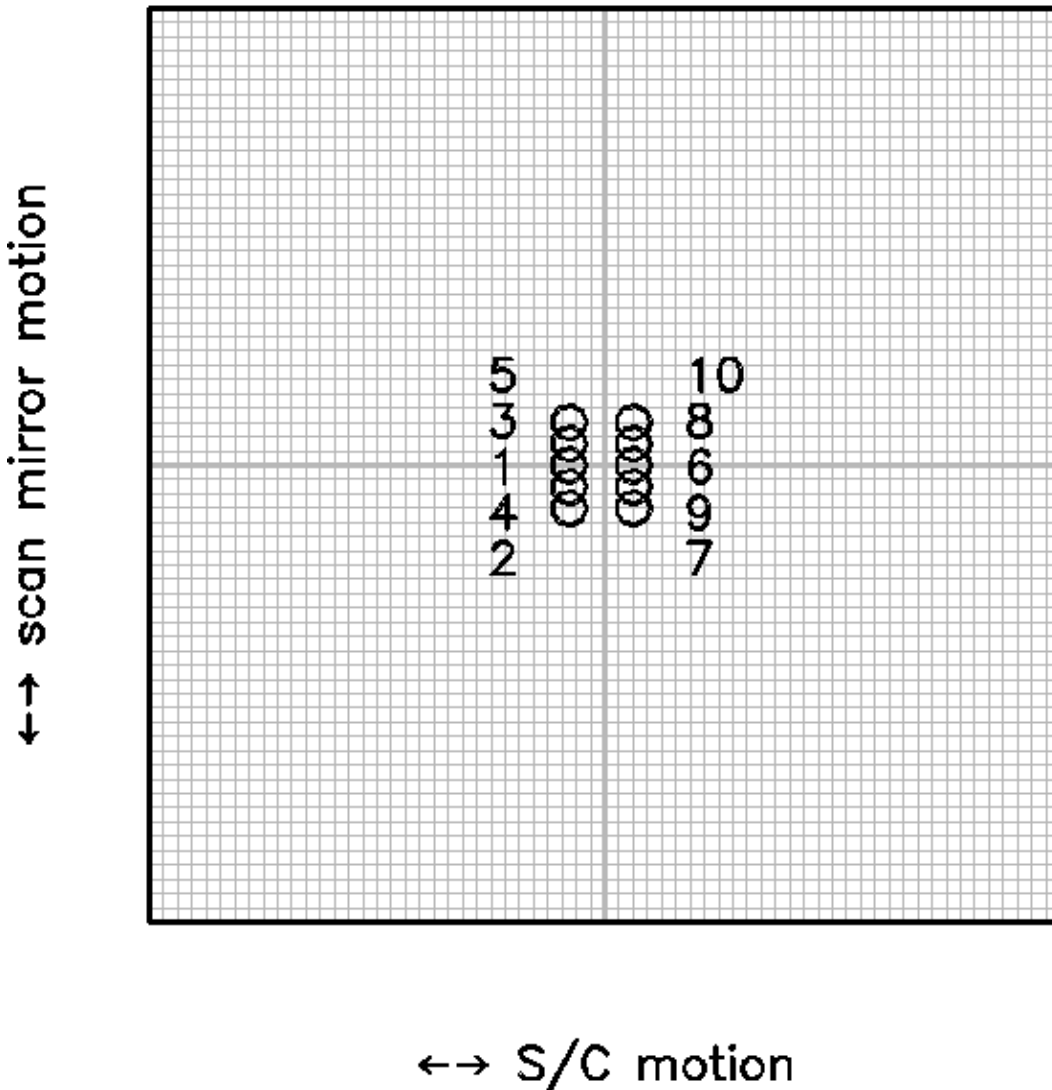
Figure 8.19: Observing a large-diameter source at 24 μm. The positions of image centers in a single cycle of a 24 μm large-source photometry observation are shown relative to the central 64×64 pixels (shown schematically by the grid) of the full 128×128 array. Dither positions are at half-integer pixel offsets in both the scan-mirror direction and the cross-scan direction. An identical set of observations is obtained at an observer-specified sky position offset by >5′ in the scan direction. Not shown are the two extra exposures that are obtained for each AOR, which are taken at positions 1 and 6.

Table 8.8: Source positions for 24 μm large source photometry (in units of pixels, the origin being the center of the array).

Frame number	X position	Y position
1	-2.5	0.0
2	-2.5	-3.1
3	-2.5	1.5
4	-2.5	-1.6
5	-2.5	3.1
6	2.0	0.0
7	2.0	-3.1
8	2.0	1.5
9	2.0	-1.6
10	2.0	3.1

As before, Table 8.8 contains the resulting approximate image positions in units of pixels, with the origin taken to be the center of the array. Array distortion and the twist of the scanning mirror with respect to the array columns have not been taken into account in this table, but the positions are good enough for observation planning purposes.

8.2.1.2.3 *Photometry – 70 μm Compact Source (Small Field), Default Pixel Scale (Coarse Scale)*

Typical photometry observations at 70 μm are made using the nominal 9.8″ pixel scale. The super-resolution mode at 70 μm also provides excellent data for photometry, but is considerably more expensive in terms of observing time overhead and is also about 4 times less sensitive. In the normal photometry observation of a compact ($< 1'$) target (see Figure 8.20), a pattern of observations is made that is similar in concept to the 24 μm compact source photometry observation. As a result of the on-orbit realities, we made changes to our pre-launch plans for this mode; experienced MIPS observers will notice that the currently implemented AOT resembles the pre-launch AOT, but shifted to the less noisy side (side A) and with the amplitudes of the offsets correspondingly scaled down in the cross-scan direction.

To start the observation sequence, the source is centered to the left of the array center by 9 pixels (98″), 1 pixel to the left of the center of side A (the good side), and the scanning mirror is chopped in a symmetrical fashion, above and below the array center to obtain a column of 6 images — the first and last are both at the central position. The spacecraft is then slewed to position the source 7.5 pixels to the left of center (74″), 0.5 pixel to the right of the center of side A, and another column of images is obtained. The resulting approximate image positions are given in Table 8.9, where the coordinates are in pixels and the origin is taken to be the center of the array. Array distortion and twist have not been taken into

account in this table. In reality, the images will not fall exactly on the positions given in this table, but these positions are good enough for observation planning purposes.
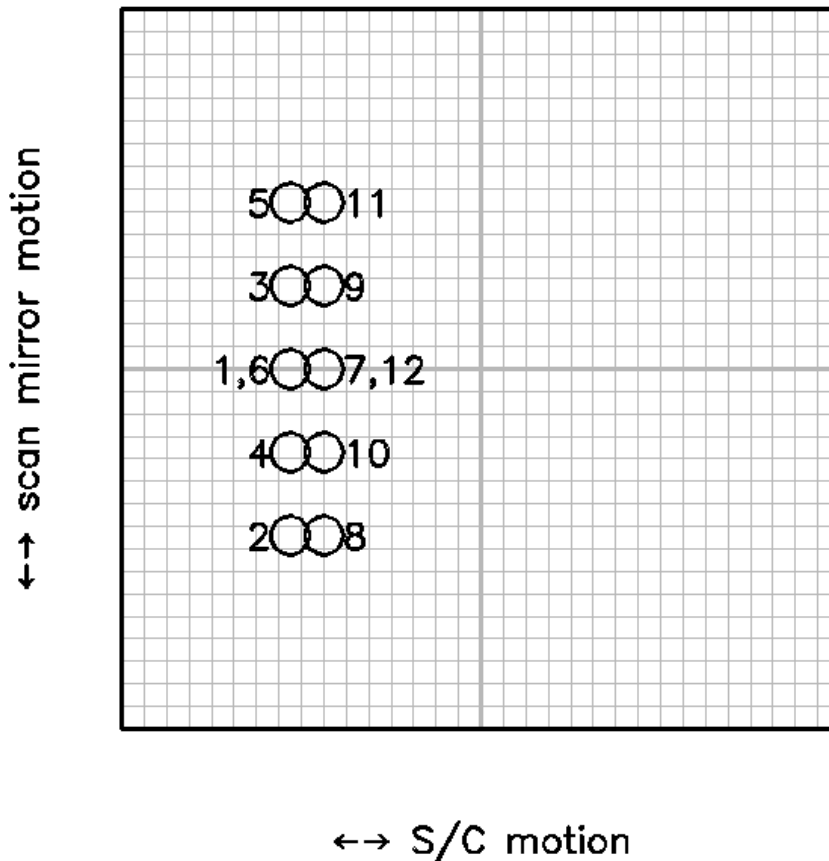
Figure 8.20: Photometry of a compact source with the 70 μm array. The positions of the target in each of the 12 frames obtained in the first cycle of the standard photometry observation at 70 μm are shown by the circles, which are the size of the PSF at the FWHM. The entire 32×32 array is shown with pixels represented schematically by the grid. The dither pattern involves half-integer pixel offsets in both directions. Repeat cycles of the photometry AOT will provide just images 2–6 and 8–12, not all 12 shown here.

A stimulator flash frame follows the first and the sixth exposures in the sequence. A telescope nod then moves the target (see the figure) and the pattern of frames and stim flashes is repeated, completing a single cycle of the AOT and producing 10 target images, with an extra pair of images on the first cycle (and perhaps a few later cycles as well). The observer specifies the number of times to repeat the complete cycle; if multiple cycles are requested, several sets of frames may be taken at one nod position before switching to the other position. The observer has no control over the relative placement of the stim flashes or the frames. The relative positions of these images are chosen to provide ½ pixel sampling of the

PSF in both the scan and cross-scan directions. The positioning of the target on the array is illustrated in Figure 8.20 and Figure 8.21; in the first of these figures, the circles show the three consecutive pairs of positions for the image with their diameters being the FWHM of the Airy disk.
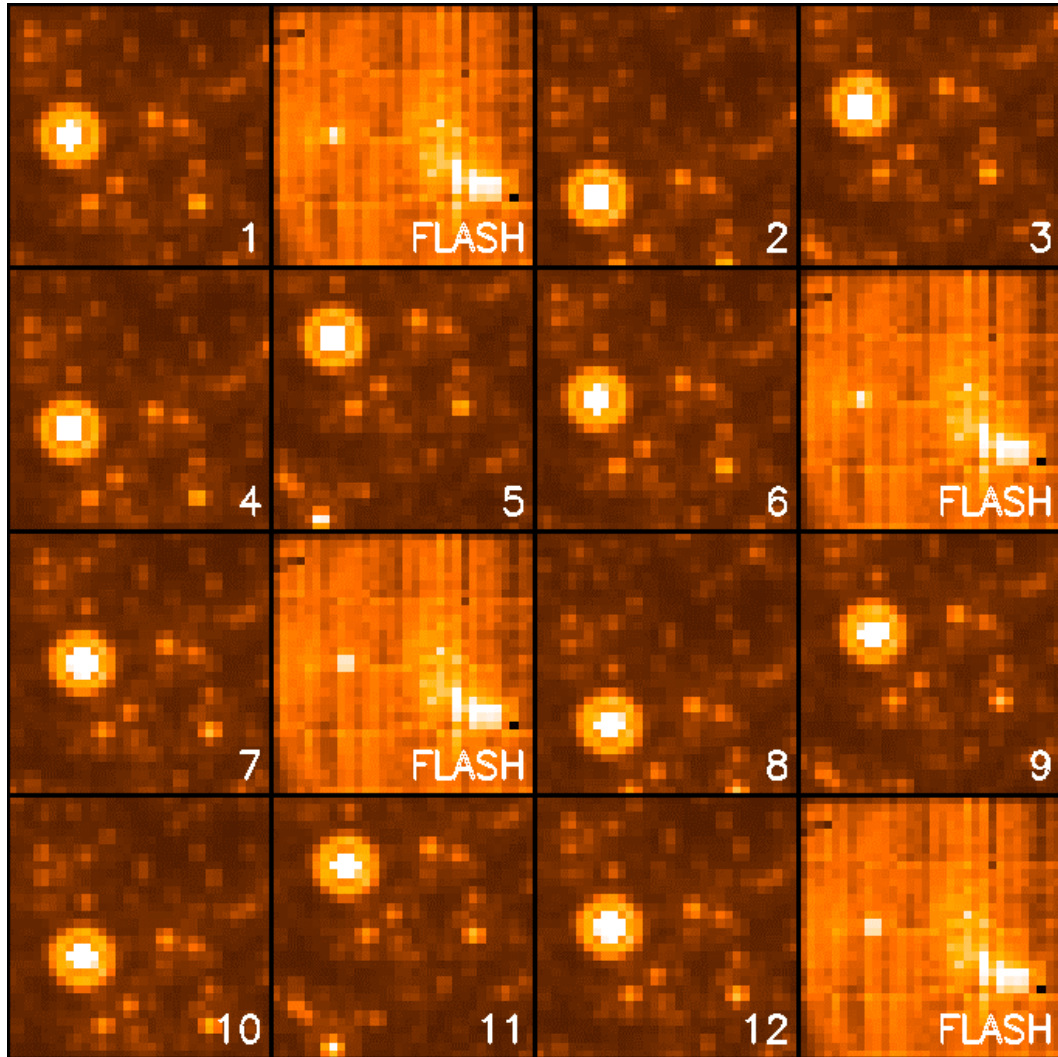
Figure 8.21: Simulated source detections on the array during the compact source photometry AOT in Figure 8.20. Time runs across the top row from left to right, then across the second row from left to right, and so forth. The observation starts with a source frame and stimulator flash, then several scan-mirror dithered frames. The sequence is completed with another stimulator flash. The spacecraft is then offset, and the above sequence of dithers and images is repeated. In repeat cycles of the photometry AOT frames 1 and 7 *are omitted*, so the AOT basically provides 10 source images per cycle, not all 12 shown here. Note that the source stays on side A of the array.

These photometry data provide multiple independent images with sampling of various positions on the pixels and a level of oversampling that is useful in extracting diffraction-limited images. Note that careful combination is required to gain all the benefits of this sampling, although for simple photometry, integer pixel shifts are adequate. Furthermore, the scan mirror motion is not perfectly

parallel to the columns of the array. Nonetheless, the pattern provides a redundant level of oversampling that is useful in extracting diffraction-limited images.

Table 8.9: Source positions for 70 μm default pixel scale compact source photometry (in units of pixels, the origin being the center of the array).

Frame number	X offset	Y offset
1	-9.0	0.0
2	-9.0	-7.4
3	-9.0	3.7
4	-9.0	-3.7
5	-9.0	7.4
6	-9.0	0.0
7	-7.5	0.0
8	-7.5	-7.4
9	-7.5	3.7
10	-7.5	-3.7
11	-7.5	7.4
12	-7.5	0.0

Table 8.10: Source positions for 70 μm default pixel scale large field photometry (in units of pixels, the origin being the center of the array).

Frame number	X offset	Y offset
1	-9.5	10
2	-9.5	-2
3	-9.5	12.5
4	-9.5	0.5
5	-9.5	15
6	-9.5	3
7	-7	-8.5
8	-7	3.5
9	-7	-11.5
10	-7	1
11	-7	-14
12	-7	-1.5

Given the characteristics of the MIPS detectors, it is best to build up long integrations with long exposure times, e.g., large total integration times built from 3 s exposures will not produce as high-quality a final product as fewer cycles of 10 s exposures.

8.2.1.2.4 Photometry – 70 µm Large Source (Large Field), Default Pixel Scale (Coarse Scale)

For 70 µm imaging of a source up to ~ 2′ in diameter, the limited motion possible with the scan mirror requires that the source be referenced to sky in two "halves," as illustrated in Figure 8.22. As above, as a result of the on-orbit realities we made changes to our pre-launch plans for this mode; experienced MIPS observers will notice that the currently implemented AOT resembles the pre-launch AOT, but shifted to side A, and with the amplitudes of the offsets correspondingly scaled down in the cross-scan direction.
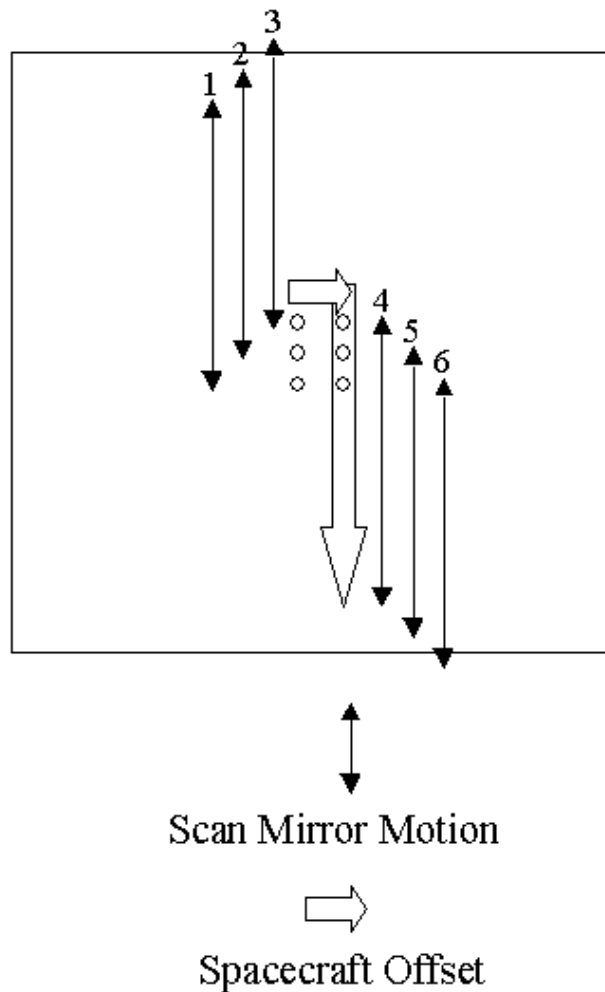
Figure 8.22: Photometry of a large source at 70 µm. The source is observed such that the coordinates are placed near the array center in the pattern shown. Spacecraft and scan mirror motions are made in such a way that the sky background is observed in two halves on both sides of the image in the in-scan direction. Note that for simplicity, the array is portrayed as a square; in-orbit realities mean that only half of the array is functional, and the pattern as portrayed here is centered on side A.

The scan mirror starts near one end of its travel within the large-scale 70 μm range, and the spacecraft is pointed to place the source 2.5 pixels below the center of the array and 0.5 pixel to one side of the center of side A. The scan mirror is set to have a throw of 2′, and the first pair of exposures is taken at the starting point (target image) and 2′ away in-scan (sky image). A second target – sky pair is obtained with the scan mirror advanced 2.5 pixels (24.6″), and a third pair obtained after advancing another 2.5 pixels. Next, the spacecraft is maneuvered 1.25 pixels in cross-scan and 2′ in-scan, and the scan mirror is then used to obtain another three pairs of exposures, starting with a sky image, then chopping back to get the target. The overall sky coverage is illustrated in Figure 8.23.
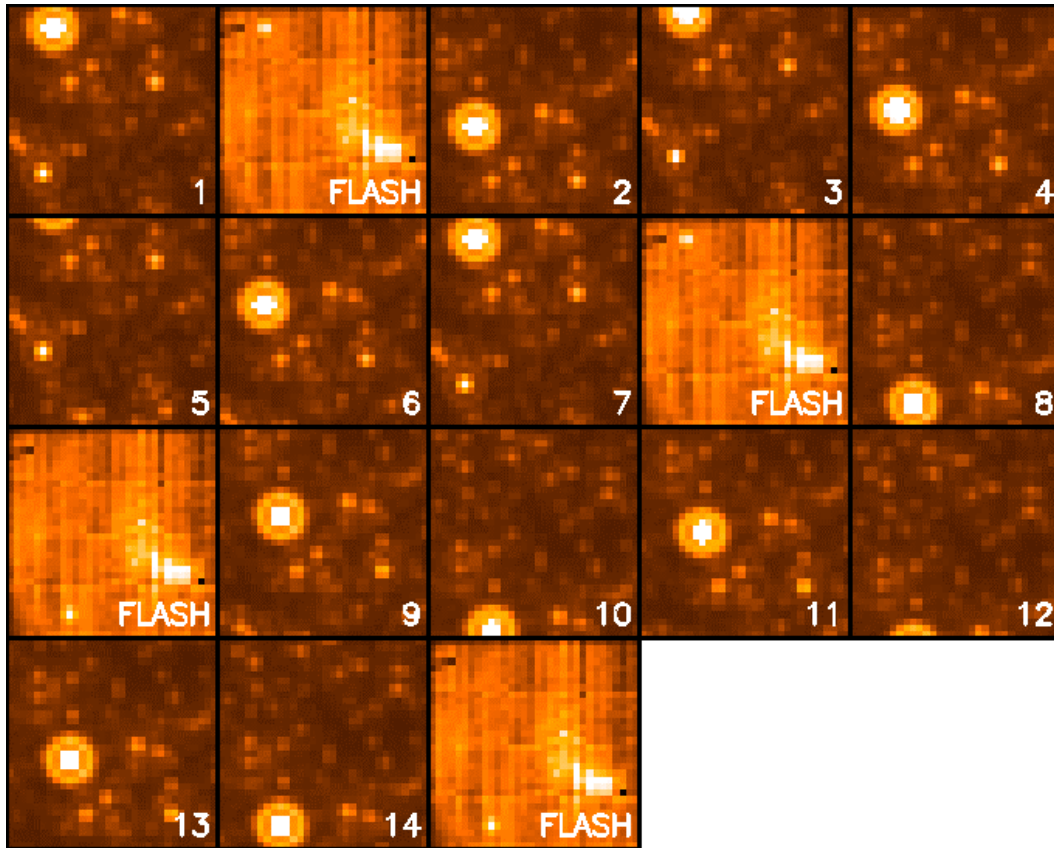
Figure 8.23: Simulated source detections on the array during the large source photometry AOT shown in Figure 8.22. Time runs across the top row from left to right, then across the second row from left to right, and so forth. The observations start with an image in the sky position, then a stimulator flash, then the first source image. Three on source-off source pairs of images are taken, then a final stimulator flash. The spacecraft is then offset, there is a stimulator flash, three more on-off image pairs are taken, and the cycle is completed with a final stimulator flash. In repeat cycles of the AOT, frames 1 and 7 *are omitted*, so the AOT basically provides 6 source-sky image pairs per cycle. Note that the source stays on side A of the array.

Note that the viable part of the array is a 2.5′×5′ rectangle, but since the orientation of the field of view on the sky is a function of time and generally observers cannot specify the time of their observation (see section 3.2.5), the largest object that should be observed using this mode is <2′.

Because of the smaller effective size of the 70 μm array compared to previous expectations, some users will need to efficiently map the approximately 5′×5′ (2.5′×2.5′ fine scale) area covered by the original large-field AOT. In order to do this, we suggest that the user create a cluster target with offsets of (0, +80″), (0, -80″) in array coordinates observing "offsets only." This will allow recovery of the full area of the original AOT while saving the slew tax from mapping with 2 AORs.
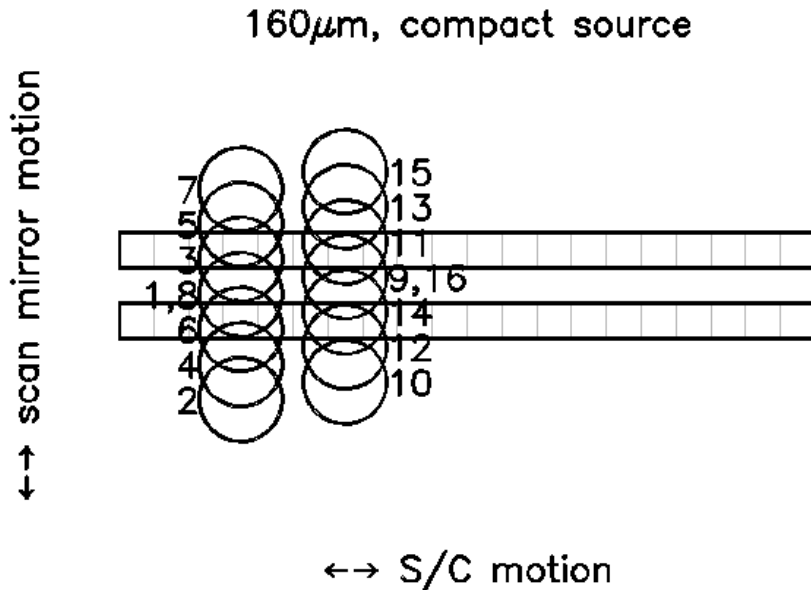
Figure 8.24: Positions of the source relative to the 160 μm array during the compact source photometry AOT. The scan mirror positions and spacecraft offset (between frames 8 and 9) provide sampling of the image at ½ pixel offsets in both dimensions. *A single cycle of the AOT provides only 2 complete images of the source because the array contains only 2 rows of pixels.* **Observers should request a minimum of 2 cycles to obtain at least minimal data redundancy. Frames 1 and 9 are omitted from repeat cycles of the AOT.**

8.2.1.2.5 Photometry – 160 μm, Compact Source (Small Field)

At 160 μm, the array width is less than the FWHM of the Airy pattern and is unfilled (2×20 pixel rows with a one-pixel-wide blank row between), so a series of frames must be taken to build up an image of the source. As a result of the on-orbit realities, we made changes to our pre-launch plans for this mode; experienced MIPS observers will notice that the currently implemented AOT resembles the pre-launch AOT, but shifted to the side away from the inoperative

readout and with the amplitudes of the offsets correspondingly scaled down in the cross-scan direction.

The sequence currently implemented is illustrated in Figure 8.24. At the beginning of the observation, the scan mirror is near the center of its range, and the telescope is used to center the source halfway between the two rows of the array and offset to one side in the cross-scan direction (pixel position ([−6.5, 0.0]). An image is taken, and is followed by a stimulator flash image. The scan mirror is then used to build up a filled image by moving the source to pixel positions [−6.5, −2.0], [−6.5, +1.0], [−6.5, −1.0], and [−6.5, 2.0]. A final image is taken at the starting position, and is followed by a stimulator flash image. The telescope then nods to place the source on the other half of the best side of the array, and offset by ½ pixel in the in-scan direction (pixel position [−3.5, 0.5]). The above sequence of images and scan mirror motions is then repeated, completing one cycle of the AOT. The relative positions of these images are chosen to provide ½ pixel sampling of the PSF in both the scan and cross-scan directions. If multiple AOT cycles are needed, several sets of frames will be taken at one nod position before switching to the other position – this is transparent to the observer. The basic AOT cycle provides only two complete images of the target when all dither positions are mosaicked together, with an extra pair of sub-images on the first (and perhaps a few later) cycles.
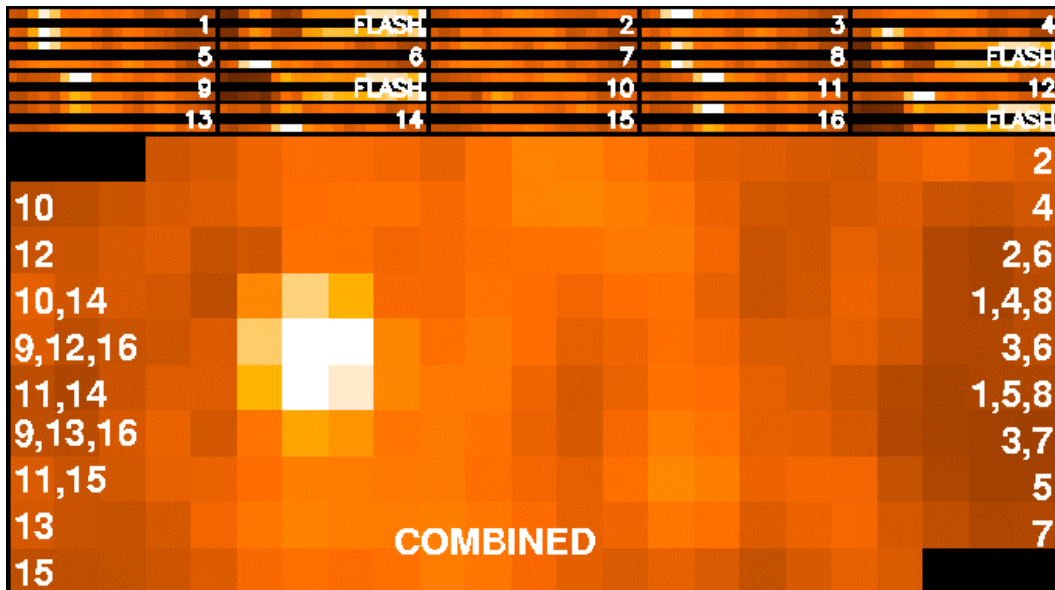
Figure 8.25: Simulated visualization of individual frames in the first cycle of the 160 µm small source photometry AOT shown in Figure 8.24 (upper panels), and a full mosaic of those frames (lower panel) with the 2 source images coadded. Note that pixels in the lower panel are shown twice as large as they are in the upper panels. Numbering along the sides shows the positions of the rows of the array during the execution of the AOT. Note that the source is offset in the final mosaic as a result of the AOT modifications to accommodate the dead readout (not shown).

It is strongly recommended that observers specify at least two cycles of the 160 µm small field photometry AOT (i.e., 4 complete images) in order to provide the minimum level of redundancy. Observers are also encouraged to use the new enhanced mode (see section 8.2.1.2.7), particularly for point source photometry.

Figure 8.25 illustrates the individual pointings for a complete cycle, and a coadded mosaic (with the two individual images coadded in the central region of the mosaic). As can be seen in the figure, the AOT provides in the end an image of the target that is offset, i.e., more of the background sky is observed on one side of the target than the other.
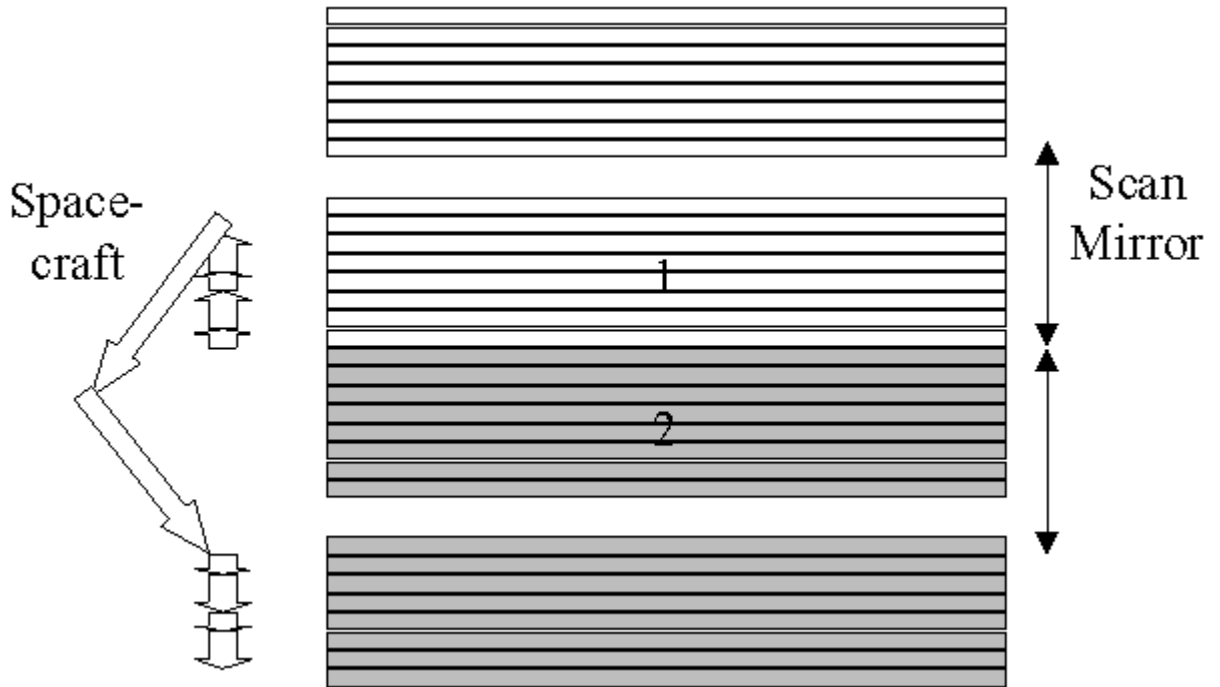
Figure 8.26: Photometry of a 4′×5.3′ region at 160 µm. A first sequence of measurements uses the scan mirror to offset the field by 2.5′ and the spacecraft to make smaller offsets to provide a filled 2′×5.3′ photometry region with a similar sky region. Compare with Figure 8.25. The spacecraft is then maneuvered to provide a second set of measurements on the opposite side of the source.

8.2.1.2.6   Photometry – 160 µm, Large Source (Large Field)

If the source to be measured is large, then a procedure similar to the 70 µm large-source one is used: first a 2′×5′ target field and corresponding sky field are imaged, then a telescope offset allows imaging of a second target field abutting the first, along with an additional sky field. This results in a fully sampled image over a ~ 4′×5′ field, as illustrated in Figure 8.26. This large-source AOT cycle at 160 µm provides only one fully sampled image of the source. **In order to obtain good results, observers are strongly encouraged to specify at least four cycles**

of the large-source 160 μm AOT, so that redundancy can be used to reject artifacts in the images. Figure 8.27 shows how the large-source AOT can be used to image a moderately large target, and the placement of the individual images relative to such a target. If multiple cycles of the AOT are requested by the observer, they are automatically dithered to provide sampling on ½ pixel centers. Stimulator flashes are inserted at the beginning, middle and end of the sequence to calibrate the data, but are not shown in the figures.
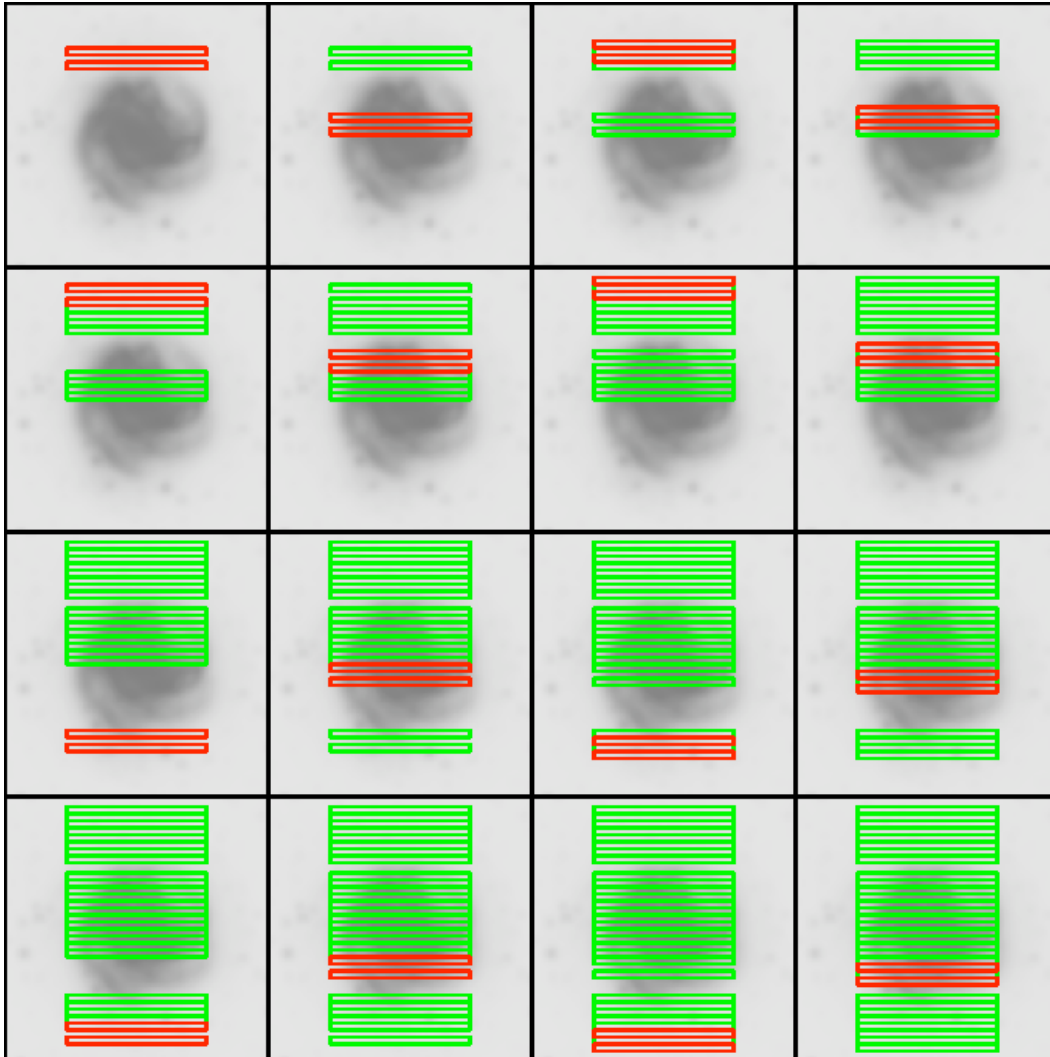
Figure 8.27: The data acquisition sequence for the 160 μm large-source option. See also Figure 8.26. The image and sky data are acquired by a series of mirror motions (source to sky) and spacecraft motions to fill in the field of view. The sequence goes from left to right and top to bottom. The current array field of view is shown in red (or bold) and already acquired fields are shown in green (or light).

8.2.1.2.7   Photometry – 160 μm, Enhanced (Small Field)

One of the conclusions from analysis of nearly three years of 160 μm calibration data, plus that of some technically challenging science programs (e.g., new planets in the Solar System or Kuiper Belt objects), is that to improve both the

photometric accuracy and repeatability of the 160 μm small field photometric observations, the AOT needed to be modified. This has led to the development of an "enhanced" 160 μm small field photometric mode by the Instrument Team and SSC staff, to take full advantage of the 160 μm array capabilities.
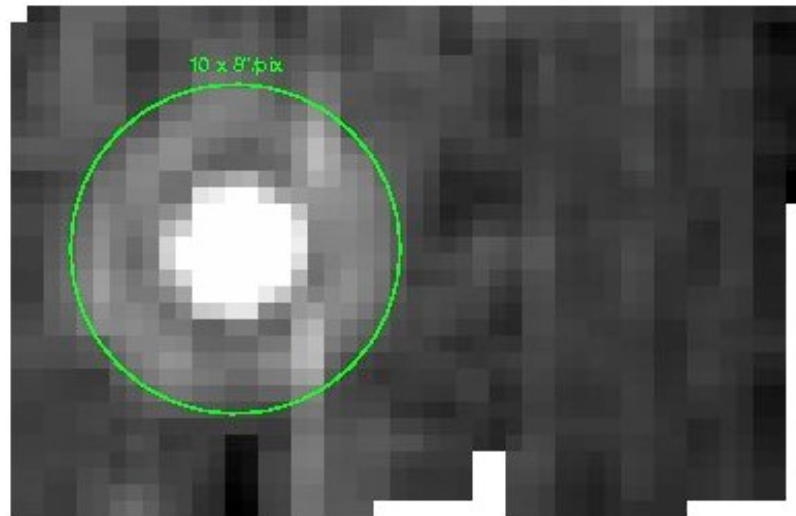
Figure 8.28: Asteroid Harmonia observed using the 160 μm enhanced mode during MC42 (July 2007). This asteroid is bright (approximately 3 Jy) and samples the MIPS 160 μm upper limit point source calibration (see Stansberry et al. 2007). The enhanced mode FOV clearly encloses the first Airy ring (the green circle is 1.5′), allowing for a good PSF fitting and/or a better aperture photometry with background subtraction.
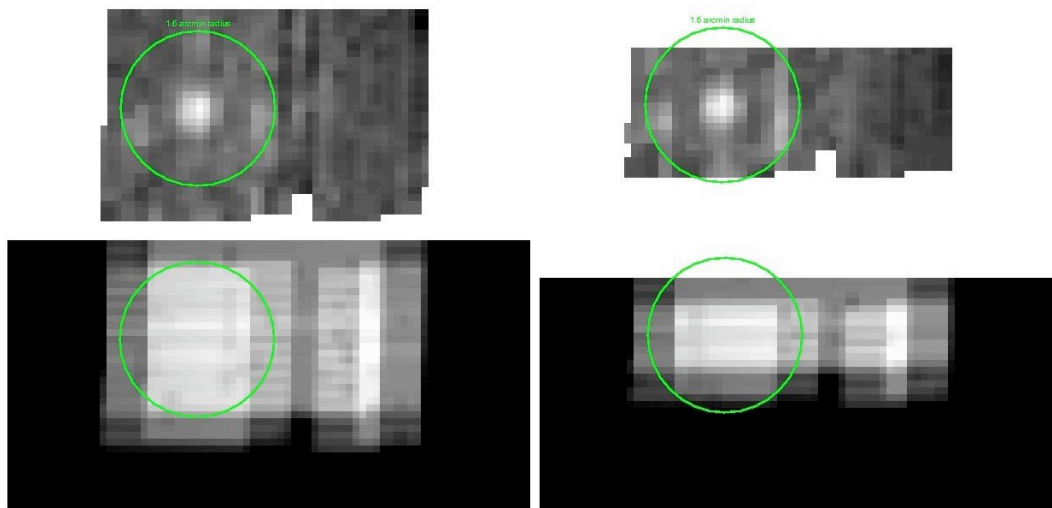
Figure 8.29: A comparison of photometric observations of 3C371 at 160 μm using the enhanced mode (left) and the "default" small-field mode (right), obtained during MC41 (2007 June). The bottom panels are the corresponding coverage maps. The observations utilized 4 cycles and are displayed to the same greyscale level. The green circle corresponds to 1.5′ and encloses the 160 μm first Airy ring (see Figure 8.28). The mean coverage is about the same (10% higher for the small-field mode); however, to cover a similar FOV the small field would need twice as much time as the enhanced mode, and the coverage would not be as uniform.
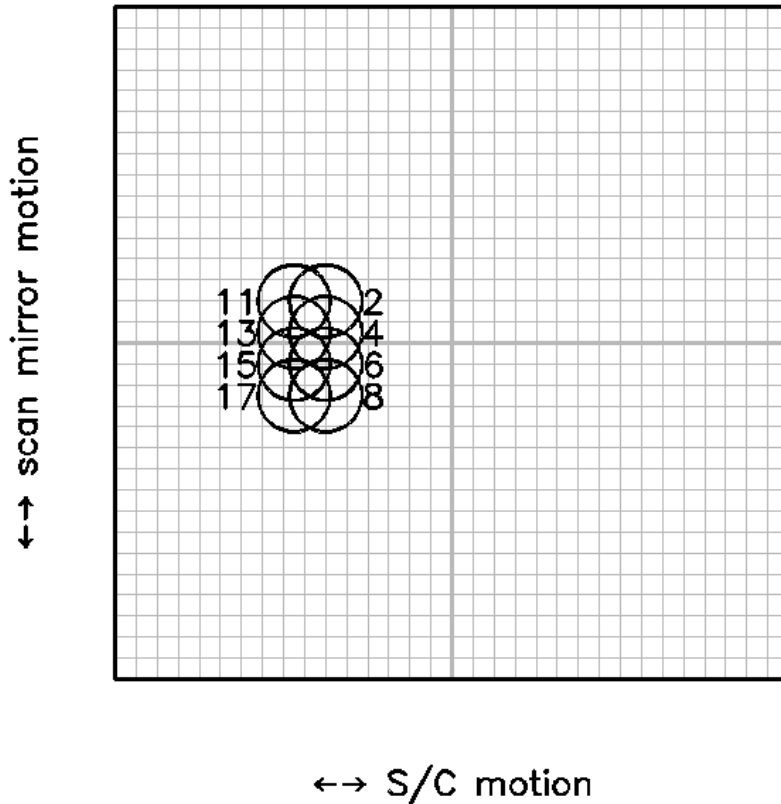
Figure 8.30: Source positions on the 70 μm array during the super resolution AOT for a
compact source. These frames cover the source region; every other frame is of a sky position
roughly 4′ away, and are obtained using scan mirror deflections. The pixel scale here is 5″,
so the circles representing the FWHM size of the PSF are correspondingly larger than for
figures illustrating observations using the nominal (9.98″) pixel scale at 70 μm.

The new 160 μm "enhanced" photometric mode relies on the same principles of
the small field photometry, but provides a larger field of view and a more uniform
coverage at a given single nodding position. This is accomplished by increasing
the number of DCEs, modifying the stim cycle, and optimizing the scan mirror
dither pattern. Preliminary tests during commissioning of this mode, along with
more recent calibration observations (in 2007 June and July) confirm an
improvement in sensitivity of 15%. **This mode represents a significant
improvement for point source photometry at 160 μm and should be the mode
observers employ.** Figure 8.28 and Figure 8.29 show a sample of observations
obtained using the 160 μm "enhanced" mode of two red sources, the asteroid
Harmonia and the extragalactic source 3C371, both part of the MIPS 160 μm
calibration plan. Observers can propose to use this mode during Cycle-5.

The 160 μm "enhanced" photometric mode takes 30 DCEs per cycle, in
comparison with the 20 DCEs obtained by the small field photometry mode. The

input parameters for this mode can be accessed through Spot. Table 8.16 includes some of the expected integration times and overheads for one and eight cycles.

8.2.1.2.8 Super Resolution Mode

In order to acquire data suitable for super resolution post-processing, we take data in a manner that provides a high level of spatial sampling of the Airy pattern and redundant imaging at all three wavelengths. Simulations (Bippert-Plymate, Rieke, & Paul 1992; Astronomical Data Analysis Software and Systems I, A.S.P. Conf. Series, Vol. 25, D. Worrall, C. Biemesderfer and J. Barnes, eds., p. 205) show that such oversampling enhances the ability of computer processing to improve the angular resolution that is achieved in final data products. The standard 24 μm photometry observation described above is suitable for super resolution because of the $\sim \lambda / 2.5D$ pixels and highly redundant images. Such is not the case at 70 and 160 μm; the nominal 70 μm pixel size is too large, and the 160 μm observation provides only 2 images of a source, which is inadequate for good image reconstruction. Here we describe procedures for obtaining good super resolution observations in these two bands. Table 8.11 provides a summary of this mode in all bands.

8.2.1.2.9 Super-Resolution – 70 μm Compact Source (Small Field), Fine Pixel Scale (Narrow Field)

At 70 μm, the default (coarse) pixel scale under-samples the Airy pattern. However, a separate "narrow field" optical train provides imaging at ½ the coarse pixel scale, and is specifically designed for studies requiring high angular resolution. Presumably the source of interest is at least potentially extended, so separate sky images are obtained by moving the source on and off the array with the scan mirror. As a result of the on-orbit realities, we made changes to our pre-launch plans for this mode; experienced MIPS observers will notice that the currently implemented AOT resembles the pre-launch AOT, but shifted to the less noisy side (side A) and with the amplitudes of the offsets correspondingly scaled down in the cross-scan direction.

To begin, the scan mirror is offset to one of the ends of motion permitted for the high-resolution optical train. Defining the center of the array as [0, 0], the telescope is pointed to place the source on pixel [*x,y*] = [–6, +2.5] and an exposure obtained. The scan mirror is moved to the other extreme of the permitted range and a sky exposure made. Next, the scan mirror is moved to pixel [–6, +0.5] and an exposure obtained, followed by a sky exposure, followed by one at [–6, –1], followed by sky, by one at [–6, 2.5], and by a sky. The telescope is then repointed in the cross-scan direction, and the sequence repeated with on-source exposures at positions [–7.5, +2.5]; [–7.5, +0.5]; [–7.5, –1]; and [–7.5, –2.5]. The actual source positions may differ somewhat from those given above and in Figure 8.30, which illustrates this mode of obtaining data; Figure 8.31 shows a visualization of the data frames relative to a hypothetical source. This sequence gives two independent images sampled on half-pixel centers. The 1.5 pixel spacing has been adopted to provide some resilience to any dead pixels, which should then not

remove two adjacent pieces of data. Additional data to increase the net integration are obtained from a slightly different starting position to dilute further the effects of any dead pixels on the final image. The observer need only select the number of basic observation cycles required. The SSC has optimized the procedure for obtaining them to achieve the best efficiency and data quality, and may do so again during the mission.
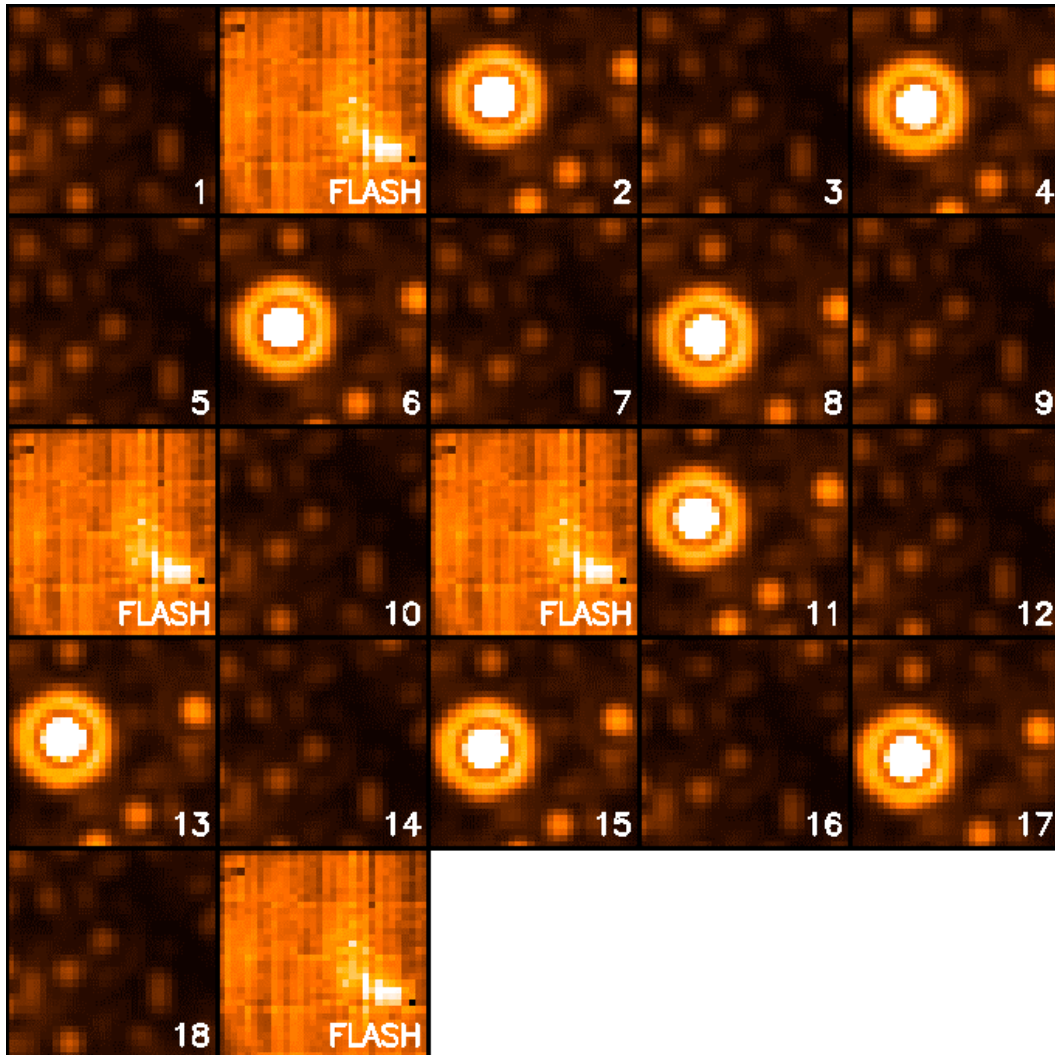
Figure 8.31: Source detections on the array during the 70 μm compact source super resolution photometry AOT shown in Figure 8.30. Time runs from left to right, top to bottom. The observations start with an image in the sky position, then a stimulator flash, then the first source image. Four on source – off source pairs of images are taken, then a final stimulator flash. The spacecraft is then offset, there is a stimulator flash, four more on – off image pairs are taken, and the cycle is completed with a final stimulator flash. Repeat cycles of the AOT omit frames 1 and 10, but still provide 8 source – sky image pairs per cycle.

8.2.1.2.10 Super Resolution – 70 μm Large Source (Large Field), Fine Pixel Scale (Narrow Field)

The 70 μm large field fine scale mode is nearly equivalent to doing a 2×2 raster map of the 70 μm small field fine scale mode which is centered on the target. It differs from the simple raster in that the scan mirror chops in opposite directions for the top and bottom rows of the 2×2 map. Thus, the result is a 2×4 fine scale map roughly centered on the target. The actual map center is offset slightly so that the central target position is not on the edge of the array in at least one of the raster positions. See Figure 8.32. Note that the viable part of the array is a rectangle (1.25′×2.5′), but since the orientation of the field of view on the sky is a function of time and generally observers cannot specify the time of their observation (see section 3.2.5), the largest object that should be observed using this mode is <2′.
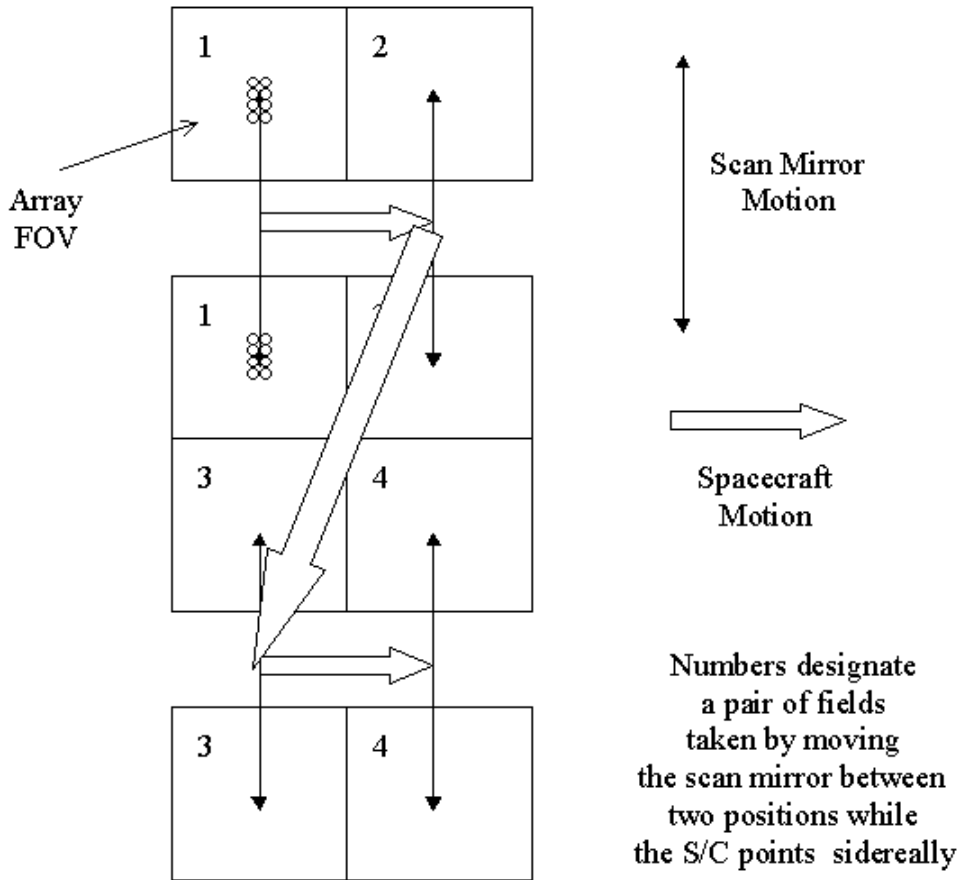
Figure 8.32: Strategy for obtaining super resolution observations over a field of 2.5′×2.5′. At each pair of fields, a set of 16 exposures is obtained similar to the pattern illustrated in Figure 8.30. Note that for simplicity, the array is portrayed as a square; in-orbit realities mean that only half of the array is functional, and the pattern as portrayed here is centered on side A.

In order to cover roughly a 2.5′×2.5′ field around the target, a cluster target should be constructed with offsets of (0, +60″) and (0, -60″) (offsets only). This saves the slew tax from mapping with 2 AORs.

8.2.1.2.11   Super Resolution – 160 µm Compact Source (Small Field)

The 160 µm pixels are nominally small enough to provide good sampling of the PSF for high-resolution image reconstruction. However, the impact of cosmic rays and the complicated behavior of the Ge detectors mean that integration time and PSF sub-sampling beyond what is provided by the normal 160 µm photometry mode are needed to obtain good super resolution results. Unlike the 70 µm array, there is no separate pixel scale or tailored AOT available for 160 µm super-resolution observations. The recommended method for specifying such an AOR using the available 160 µm pixel scale is described in section 8.2.3.3. Such an observation consists of 24 individual frames that provide three 2′×5.3′ images, each containing 2 images of the source, with ½ pixel sampling in both the scan and cross-scan direction. The third image has the same cross-scan direction sampling of the PSF as the first image, but is placed on different pixels of the array to provide extra protection against unusual behavior of the Ge:Ga pixels and intra-pixel responsivity variations.

8.2.1.2.12   Photometry Raster Map

Scan maps are very efficient, and obtain data at all 3 wavelengths at once. However, for objects that are somewhat larger than 2′ in diameter, but smaller than the 30′ minimum scan map size, the photometry raster map option can be used to construct a map in one or more of the MIPS bands. The choice of photometry raster maps vs. scan maps for any given observation can be made via empirical comparisons using Spot visualizations and observing time calculations. Here we discuss the raster map option in more detail.

The raster map uses the small field dither pattern identical to a single point Photometry/Super Resolution observation (same scan mirror and spacecraft offsets as described earlier). After completing observations in this fashion in a single field-of-view, the spacecraft is offset to an adjacent observer-specified field, and the normal photometry dither pattern is repeated. The observer specifies the number of nearly-full-frame offsets in the along scan and cross-scan directions. This limited raster mapping option may be expanded in complexity in a phased way if the need is apparent later in the mission. The raster option should return high-quality 24 µm results because of the high stability of that array. Observers should specify offsets that give some overlap between fields if they are interested in extended emission within the field of their raster map.

As a result of the on-orbit realities, the positioning of the target in 70 µm raster maps changed (as for regular 70 µm photometry). The Spot visualizations reflect this change.

For observers planning raster maps at 24 μm, for full coverage, observers should use at most ½ array offsets.

Another solution to this general problem of covering a small region of sky is to use cluster mode targets instead of raster maps; see section 8.2.3 below or the Spot User's Guide for more information about cluster mode targets.

Observers should use always Spot to visualize observations and be sure that the observation is doing what was intended.

8.2.1.3 Photometry and Super-Resolution Summary

Table 8.11 summarizes the expected frames per observation cycle that are obtained in the MIPS photometry and super–resolution observing modes, and the integration time per pixel for those modes.

Table 8.11: MIPS Photometry and Super-Resolution Summary

Mode	Band (μm)	Frames /Obs. Cycle[18]	Approximate Integration Time per Pixel per Cycle[19]
Compact Source Photometry	24	14[20]	42, 140, 420 sec[21]
	70	10	30, 100 sec
	160	14[22]	6, 20 sec
Large Source Photometry	24	10 / 10[20]	30, 100, 300 sec[21]
	70	6 / 6	18, 60 sec
	160	10 / 10[23]	3, 10 sec
Compact Source Super Resolution	24	14[20]	42,140,420 sec[21]
	70	8 / 8	24, 80 sec
	160	42[24]	18, 60 sec
Large Source Super Resolution	24	10 / 10[20]	30, 100, 300 sec[21]
	70	32 / 32[25]	24, 80 sec
	160	N/A	N/A

---

[18] Two values indicate number of frames on-source/off-source.

[19] For 3 and 10 second exposure times (and 30 seconds at 24 μm) respectively. Times are per pixel on a given sky position in "MIPS seconds." Actual exposure times are 1.05 times longer (see also Table 8.10). Note that Spot reports real seconds, not MIPS seconds.

[20] At 24 μm, 2 additional frames are taken per AOR, so total integration time will be longer than shown here by 2 times the exposure time. See also next note.

[21] For the first cycle in an observation at 24 μm, exposure time is 1 second shorter than shown in this table. See also previous note.

[22] The 10 160 μm frames combine to provide a 2'× >5' filled field of view containing two images of the source.

[23] The 10 160 μm frames combine to provide a 4'×5' filled field of view containing a single image of the source

[24] The 3×10 160 μm frames combine to provide a 2'× >5' filled field of view containing six images of the source sampled at sub-pixel shifts. See section 8.2.1.2.11.

[25] This indicates total number of frames, not total number of frames on-source (which is 8, just like the compact source).

8.2.1.4 Spectral Energy Distribution (SED) Mode

The Spectral Energy Distribution (SED) mode applies only to the 32×32 Ge:Ga array, since it requires an offset of the scan mirror that deflects light away from the optical trains for the other arrays. This operating mode provides long-slit, low-resolution (R ≈ 15 – 20) spectra in the far infrared (52 – 97 μm), with a dispersion of 1.70 μm/pix via a reflection grating. The SED optical train illuminates a slit approximately two detector pixels (~20") wide and 12 pixels (~2') long where the full wavelength coverage is available. An inoperative detector module at the blue end of the dispersion of side A restricts the wavelength coverage to only 65 – 97 μm over the last 4 columns of the 16x32 detector array. Only side A data is calibrated for the SED mode.
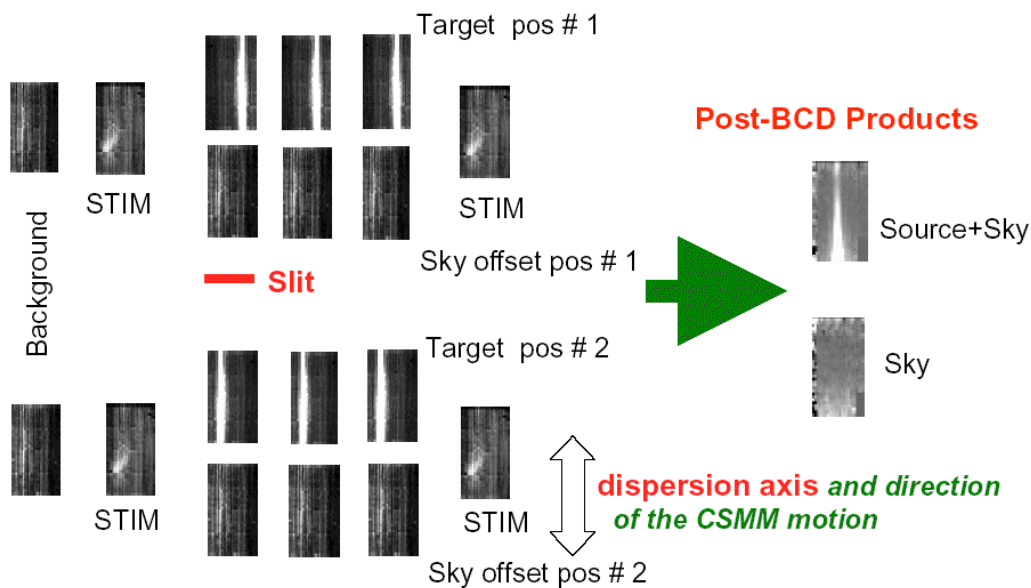
Figure 8.33: Schematic representation of SED mode; see text. These are real data.

SED observations provide pairs of data frames with the target in the slit and at a nearby off position. The scan mirror is used to chop between the target and the off position, and the observer can pick a chop throw of +3′, +2′, +1′, or -1′. An SED observation begins with an image at the off position followed by a stimulator flash frame. The scan mirror brings the target into the slit, and a series of 3 target – sky image pairs is taken with the scan mirror chopping between the two positions. The series ends with a stimulator flash frame following the final sky frame. The telescope then nods to move the target along the slit to a second position, and the above sequence of images, stimulator flashes, and scan mirror motions is repeated. If additional integration time is needed, the sequence is repeated except that the first sky – stim flash pair of frames is omitted. Consequently the AOT provides 6 target – sky image pairs per cycle; see Figure 8.33. *The observer specifies only the total number of basic observation cycles required*; the SSC has optimized the number of pairs to be taken at one spacecraft setting before offsetting to the other, as with the photometry mode.

As a result of the on-orbit realities, we made changes to the original SED AOT, such that we avoid side B and the dead part of side A as much as possible. Note that full spectral coverage is limited to only 32×12 because of the dead readout on side A; the spectral coverage for 4 columns of the array is reduced to about 65-95 μm.

8.2.1.5 Total Power (TP) Mode

MIPS is optimized to provide calibrated images of sources that are small enough that they can be chopped on and off of the arrays, particularly at 70 and 160 μm. The Total Power (TP) mode AOT provides a way to accurately measure extended emission as well, by chopping between an internal dark position and the sky for each array. Experienced MIPS observers will note that the strategy for the 70 μm array changed from pre-launch expectations; now the 70 μm TP observations use side A of the array (the good side) and do indeed chop between the sky location to be measured (using the default 9.98″ pixel scale), and an internal dark position.

The basic observation sequence begins with a frame at the dark position followed by a stimulator flash frame. Six more frames are taken at the dark position to allow the latent image from the stimulator flash to fade, then the scan mirror is used to change to the sky position, and a single image is taken of the source region. The scan mirror then moves back to the dark position, and 6 more frames are taken, allowing time for the latent image of the sky to fade. A final stimulator flash completes the cycle. This pattern (less the initial dark and stimulator flash frames) is repeated 5 times, providing 5 sky – dark image pairs per AOT cycle. This time-consuming process should provide excellent data on the germanium arrays. At 24 μm, image latents are better behaved. However, a 24 μm stimulator flash is included in this AOT (the only AOT that uses the 24 μm stimulator) in order to provide a reference signal for applying the droop correction to the data.

8.2.2 Performance of the Instrument

The performance of MIPS in actual use depends on many factors. The performance of the arrays themselves is important, but telescope and optical throughput also play an important role. **An important factor for observers to consider when predicting the sensitivity they will achieve in a given measurement is the sky background present in the field of interest.** Sources of background flux vary with wavelength and position on the sky, and must be evaluated with care when planning observations with MIPS. The character of the far-infrared background as it is understood today is described below. However, MIPS detects the background at sensitivities and spatial scales which have not previously been achieved, so users should be conservative when estimating a signal-to-noise ratio. There is a resource available for professional astronomers on the SSC website entitled "The Infrared Compendium," which includes sections on infrared backgrounds.

8.2.2.1 Sources of Noise and Confusion

At 70 μm, zodiacal light (thermal emission from warm dust within the Solar System) and Galactic cirrus (thermal emission from warm dust clouds in the Galaxy) dominate the background. Zodiacal light is smooth on the scale of a MIPS frame, and decreases by roughly a factor of 4 between sight-lines near the ecliptic plane and towards the ecliptic pole. Zodiacal emission at a particular RA and Dec also varies as a function of time due to the orbital motions of both Spitzer and the zodiacal dust. As a rule-of-thumb, the zodiacal emission can be expected to change by about 1% over a one-day period. Galactic cirrus is concentrated near the Galactic plane, and varies on many spatial scales, including scales that are not resolved by MIPS. Structures within the cirrus include wisps and small knots, and are thus of interest for planning, not only because of the background flux they contribute, but also because they can potentially be confused with sources of interest. The level of cirrus varies greatly, and must be evaluated in detail (see Estimating Background Noise, section 8.2.2.2 below) for observations where high sensitivity at 70 μm is desired.

Galactic cirrus contributes somewhat to the confusing structure of the background, but especially at long wavelengths, this structure is dominated by emission from partially-resolved and unresolved distant galaxies. Observations that are subject to this background are frequently referred to as "confusion limited," because the spatial structure resembles the superposition of many point sources, and determining whether a particular faint source is a part of the background or is a target can be difficult. The structure of the background effectively increases its noise contribution above what would be calculated based on the flux contributed by all of the sources in a region. The structure contributes directly to the variance of the background as it would be measured by aperture or PSF-fitting photometry extraction software, for example.

The accuracy of photometry at 70 and 160 μm will often be confusion-limited. Because MIPS provides much smaller effective beams and higher sensitivity than any previous mission, predicting the confusion limit set by such sources is difficult. Pre-launch estimates of the 1σ confusion limits ranged from about 0.5 to 1.3 mJy at 70 μm, and from about 7 to 19 mJy at 160 μm (Xu et al., 2001, ApJ, 562, 179; Franceschini et al., 2002, ESO conference proceedings, astro-ph/0202463; and Dole et al. 2003, ApJ, 585, 617). Early mission data from MIPS produced the following source density criterion (SDC) limits for confusion, due to *extragalactic* sources (Dole et al. 2004, ApJS, 154, 93): 56 μJy at 24 μm, 3.2 mJy at 70 μm, and 40 mJy at 160 μm. The 5σ photometric confusion limit at 70 μm is 1.5 mJy (Frayer et al. 2006, ApJ, 647, L9). At 24 and 70 μm, the confusion mainly results from the high density of resolved sources; at 160 μm, the confusion arises from a population fainter than the sensitivity limit; see Dole et al. (2004, ApJS, 154, 93), available on the SSC website. There is a range of values expected, and moreover it can be a function of where exactly one looks in the sky. Note that confusion due to *Galactic* sources is a strong function of position.

Other factors may influence the effective confusion limit for a particular observation. In some instances it may be reasonable to integrate somewhat below the level of the confusion, for example when the observer has *a priori* knowledge of a source position. On the other hand, the presence of a nearby bright source with its diffraction artifacts will increase the effective confusion limit. Moving targets offer the possibility of taking a second "shadow" observation, allowing the suppression of confusing sources by subtracting them away.

While the above values should serve as a guide for determining if a particular observing program is feasible, interested observers should consult the corresponding research papers for additional details on confusion limit determination.

Observers are warned that they need to specify AORs with enough cycles to provide adequate rejection of cosmic rays and other artifacts, even if a very short integration would nominally be adequate to reach the confusion limit.
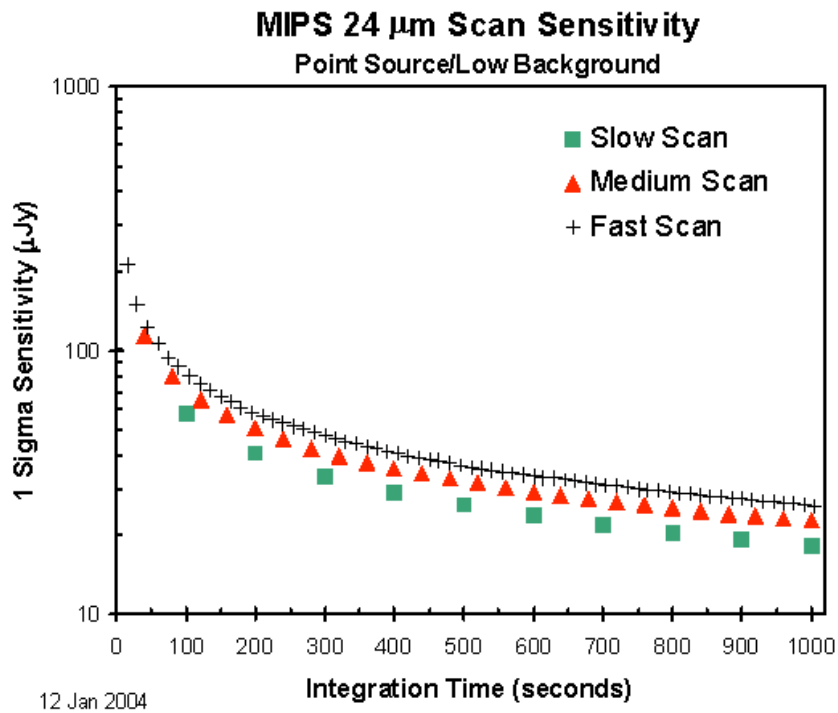
Figure 8.34: Sensitivity in scan map mode at 24 μm. At the medium and slow scan rates, the effects of cosmic rays and increased read noise largely offset the expected gains from increased exposure time.
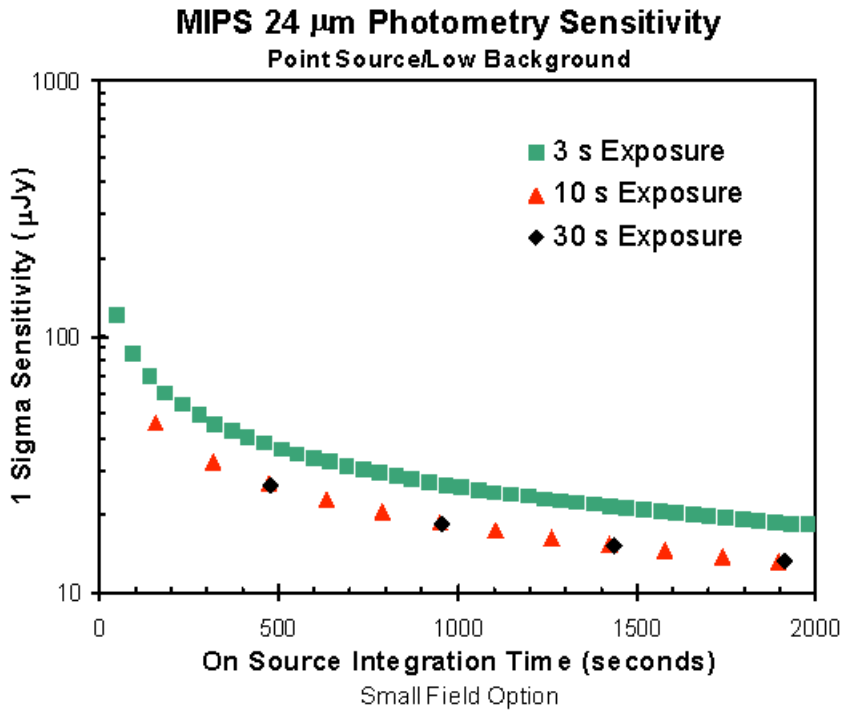
Figure 8.35: Sensitivity for photometry at 24 μm, compact source (i.e., one that can always be kept on the array during the measurement).

8.2.2.2 Estimating Background and Confusion Noise

Background noise estimation can be done using Spot. An estimation of the background and the contribution of the various components to the background (e.g., zodiacal light) can be obtained from the target selection window. Spot is designed to do this specifically for Spitzer observers, although IRSky is also available; see section 5.3.

While the noise levels predicted via Spot (or IRSky) are based on previous infrared space observations at lower sensitivity and spatial resolution than MIPS obtains, they are the best products available at this time for making such estimates and should be used for planning all observations; early results suggest that the Spot model does reasonably well (see, e.g., Dole et al., 2004, ApJS, 154, 93 and Meadows et al., 2004, ApJS, 154, 469). The only exception to this is in regions where the infrared cirrus levels are low. In these areas the IRSky/IBIS predicted noise levels may be seriously flawed because they are based on measurements with inadequate sensitivity. In these low-cirrus regions it is preferable to use a cirrus estimate based on the atomic hydrogen column density in the direction of the target. The hydrogen column density is highly correlated with the cirrus background in regions where the cirrus background is well measured. This correlation has been extrapolated to provide an estimate of low-cirrus regions. The tools for estimating cirrus in this manner are also available via the SSC

website.

Whether IRSky or Spot is used to estimate backgrounds, the aperture size used to estimate background is important, and should be clearly stated when observers describe their sensitivity calculations in proposals.

8.2.2.3 Estimating Signal-To-Noise Ratio and Integration Time

In nearly all observing circumstances, MIPS will be background-limited. That is, the dominant source of noise in the final measurement of the brightness of an object is noise contributed by emission from the background rather than the noise due to the photon statistics of the flux from the object itself. In the background-limited case, the signal-to-noise (S/N) ratio of an observation increases linearly with the flux from the object, and it increases as the square root of the integration time:

$$(S/N)_{obj} = (S/N)_0 \frac{F_{obj}}{F_0} \sqrt{\frac{I_{obj} B_0}{I_0 B_{obj}}} \qquad\qquad (8.2)$$

Where the subscript "0" refers to values obtained from the sensitivity plots (discussed below and on the SSC website), and the subscript "obj" refers to the desired object. F is the energy or photon flux from the object per second, I is the integration time in seconds, and B is the background surface brightness. (See next section for sensitivity information.)

While this equation suggests that arbitrarily high signal-to-noise can be achieved by taking long enough integrations, in practice the signal-to-noise that can be attained is finite. The details of the confusion noise sources discussed in section 8.2.2.1 above set a lower limit on both the fluxes that can be detected, and a corresponding upper limit on the signal-to-noise ratio that can be obtained for sources of low to moderate brightness. The S/N that can be achieved on bright sources is limited by the reproducibility of measurements taken with the MIPS arrays (see section 8.2.2.8). The reproducibility within a set of measurements is roughly 1% in relative flux; the repeatability from one set of measurements to another (e.g., photometry repeated on a source) is 8–10% rms. Observers are reminded that timing patterns within MIPS result in "seconds" which are 5% longer than a real second. If Table 8.11 is used to select the number of AOR cycles needed to achieve a desired integration time, the actual on-source time will be 5% greater. This is unlikely to be an issue, but might be useful in some cases where one less cycle of an AOT might be chosen, for example.

8.2.2.4 Sensitivity

Although we have just emphasized that the sensitivity of MIPS depends strongly on where one is looking on the sky, we have included here the results in the best sky regions. A program that appears feasible relative to these plots should be tested carefully using the SSC-provided tools to obtain realistic estimates of sensitivity for the background emission and local confusion limit in the region of interest. Note that starting in 2005, regular MIPS operations changes to include

"cold" (telescope at ~5.5 K) and "warm" (telescope at ~8.5 K) campaigns. The sensitivity at 24 and 70 μm remains the same for any MIPS campaign; the 160 μm array is completely saturated during MIPS-warm campaigns, so its sensitivity is of course not relevant.

The full set of plots of MIPS sensitivity for scan and photometry as a function of integration time in low, medium, and high background cases is posted on the SSC website and the sensitivities are available from within the online PET. Representative plots appear in Figure 8.34 and Figure 8.35. Note that Spot reports real seconds and the seconds used in these figures are MIPS seconds. See the SSC website for details and information on background levels used. Note also that the online SSC Performance Estimation Tool (PET) uses the latest MIPS sensitivities. In case of late-breaking updates, the SSC MIPS web pages should always be consulted for any changes to the sensitivity estimates in between SOM releases.

To guide selection of exposure times, use the same general guidelines as for photometry, e.g., 10 seconds for the Ge channels if the region is not too bright, and 3 seconds if the region will saturate in 3 - 10 seconds.

8.2.2.5 Sensitivity vs. Exposure Time

The sensitivity of the 24 μm band is shown in Figure 8.34 and Figure 8.35. The detector array is well behaved and the values for S/N can be scaled as the inverse square root of the integration time for point and extended sources. The analogous plots for the other bands are available on the SSC website. The conversion of point source sensitivity to extended source sensitivity can be found in Table 8.12.

Because of the multiple response times of the far infrared arrays, sensitivity estimation is more complex for them. The best values require high-pass filtering to remove the slow component of response; see the MIPS Data Handbook for additional information (including practical advice) on such filtering. The high-pass filtering also suppresses extended emission from sources. Good data can be obtained on such sources, but it requires careful hand processing to remove "streaks" left from slow response effects. Such removal can best be done on scan map data, and it is highly desirable to scan over the source in both directions.

Real-life specific examples of extended source work with MIPS can be found in the ApJS special issue; see, e.g., the paper by Engelbracht et al. (2004, ApJS, 154, 248) on NGC 55, Helou et al. (2004, ApJS, 154, 253) on NGC 300, and/or Regan et al. (2004, ApJS, 154, 204) on NGC 7331.

The sensitivity of the far infrared arrays can also be degraded if the stimulator flash interval is too short. In particular, the 3 sec photometry mode has a sensitivity degraded by about a factor of 2 from the numbers indicated above (section 8.2.2.1). We have left the rapid stimulator rate in this mode because it should help obtain higher photometric accuracy on bright sources. For faint sources, the 10-second photometry mode should be used in preference to building

up integration time through many 3-sec cycles.

Table 8.12: Conversion from point source sensitivity to extended source sensitivity*

Band	Conversion Formula
24 μm	Sensitivity (MJy/ster) = $1.04 \times 10^{-3} \times$ Point Source Sensitivity (micro-Jy)
70 μm default	Sensitivity (MJy/ster) = $8.81 \times 10^{-5} \times$ Point Source Sensitivity (micro-Jy)
70 μm fine scale	Sensitivity (MJy/ster) = $1.77 \times 10^{-4} \times$ Point Source Sensitivity (micro-Jy)
160 μm	Sensitivity (MJy/ster) = $2.52 \times 10^{-5} \times$ Point Source Sensitivity (micro-Jy)

*Values are for default BCD pixel scales.

8.2.2.6 Spectral Energy Distribution Sensitivity

SED mode can achieve 5σ point source sensitivities of 82, 201, and 447 mJy at 60, 75, and 90 μm in 500 seconds.   Spot allows users to specify observing parameters and calculate total elapsed time for SED mode just as for the other modes.

8.2.2.7 Saturation

In general it is undesirable to saturate the arrays by exposing them to bright sources. When they become saturated, the detector operating equilibrium is disturbed and the calibration of following exposures may be affected. These issues are particularly important with the 70 and 160 μm arrays because the readout amplifiers can no longer maintain the detector bias with a saturated signal. Consequently, there is increased cross-talk to neighboring pixels and the saturated pixel exhibits long time constant drifts that reduce its sensitivity and shift its calibration.

Nonetheless, in some circumstances, saturation cannot be avoided. The MIPS electronics provide a short-exposure look at each source along with the requested long exposure, so that information can be recovered that would otherwise be lost due to saturated signals. These short exposures provide a measure of source brightness within the first second after the array reset that begins a DCE. Measured fluxes for any source that causes saturation in <1 sec will therefore be compromised. Saturation by extended sources is a more severe problem. If many pixels attached to a single readout saturate in a given DCE, the response of the amplifier can be seriously impacted. As a result, it is important to avoid saturation on extended sources or background emission at any point within a DCE. The 1 sec saturation limits for point sources, and the 10 sec limits for extended sources, are given in Table 8.13. These limits include the effects of pixel-to-pixel responsivity variations for all three arrays, and expected responsivity changes between thermal anneals for the 70 and 160 μm arrays. Observations of targets or regions that approach these limits must be carefully

considered and planned.

Table 8.13: MIPS saturation levels

Observing Mode	Point Source Saturation in 1 sec	Extended Source Saturation in 10 sec
24 μm	4.1 Jy (*)	260 MJy/ster
70 μm default scale	23 Jy (**)	101 MJy/ster
70 μm fine scale	57 Jy	292 MJy/ster
SED (@ 60, 75, 90 μm)	250/290/1000 Jy	1087/1261/4350 MJy/ster
160 μm	3 Jy (**)	20 MJy/sr

(*) see also next table
(**) see also additional paragraphs below

Table 8.14: MIPS 24 μm point-source saturation level expressions including sky

24 μm exposure time	Saturation limit in Jy
3 sec	$1.64 - 8.0 \times 10^{-4} \times$(sky in MJy/sr)
4 sec	$1.17 - 5.7 \times 10^{-4} \times$(sky in MJy/sr)
10 sec	$0.43 - 2.1 \times 10^{-4} \times$(sky in MJy/sr)
30 sec	$0.14 - 6.8 \times 10^{-5} \times$(sky in MJy/sr)

The general case for observations involves a combination of extended and point-like emission. The implications of saturation for the 24 μm array can be computed by assuming a combination of the effects in Table 8.13, as seen in Table 8.14. For example, if the estimated background at 24 μm is 100 MJy/ster, it uses up 40% of the dynamic range in a 10 second integration. Without the background, a source of 0.5 Jy would not saturate in a 10 second integration. However, only 60% of the dynamic range is left on top of the background, so the brightest measurable source in a 10 second integration is 0.4 Jy (from the 10-sec formula in Table 8.14).

During the Cycle-1 proposal process, we received many questions regarding these two tables, specifically from observers attempting to derive Table 8.13 from Table 8.14. This is not possible! The values for truly (ideal) cases are in Table 8.13. The formulae in Table 8.14 actually describe the general case, which includes a combination of point and background flux, so they break down in the limiting cases. Further complicating matters, Table 8.14 also has incorporated into it subframes that are part of the observing sequence AND the impact of some substitution of values for point sources – there is a half-second data frame taken at the beginning of the observing sequence, and we can use the first difference to substitute in values for *point sources* that are saturated later in the exposure. This doesn't work for extended sources, as far as we know from our testing to date. The values in Table 8.13 are meant to be single-point values, with nothing hidden. The values in Table 8.14 are a better approximation to reality: a mixture of point and background sources, subframes that are part of the observing sequence, and

the use of some of these subframes to substitute saturated pixels.

For extended source 24 μm saturation limits, the 10-second limit from Table 8.13 (260 MJy/sr) can be scaled by a factor of 9.5 /(exposure time – 0.5).

At 70 μm, some targets slightly brighter than these limits can be usefully observed. However, the consequences for the immediately following reads will be significant, such as latent images and a degradation of linearity. Such observations will need to be manually scheduled near a thermal anneal, and this should be described and justified in the observing proposal. Note that, since our calibrators are not this bright, observations of objects this bright may not be as well-calibrated as fainter objects.

At 160 μm, the saturation limit for a 3 second integration is about 1 Jy. For sources brighter than this level, up to about 4 Jy, useful data will be obtained on the first few reads, but the brightest pixels will saturate before the end of the integration. As a result, there will be some degradation of the results in the readouts immediately following, but the recovery will be relatively fast.

Note that there is no equivalent to Table 8.14 for 70 and 160 μm; the equivalent effects (ramp fitting to just the points before saturation) are already included in the values given in Table 8.13 above.

Several programs had observations with severely compromised signal-to-noise ratios because the observers chose long exposure times (30 sec) in regions where the background due to zodiacal light was very high.   In these cases, the slope image becomes saturated during the exposure and a large fraction of the array has had the "soft" saturated 30 sec slope image replaced with pixels from the unsaturated, but much noisier, difference image (see section 8.1.4.3). The difference image represents only about 0.5 sec of exposure time. Thus, the use of the 30 sec exposure time has instead resulted in BCD images dominated by only 0.5 sec of exposure time! A better-designed program to detect faint sources superposed on high backgrounds would make use of coadding many shorter exposures (for photometry mode, 3 or 10 sec exposures).

The SSC documentation gives the extended 24 μm saturation limit for a 30 sec exposure to be 84 MJy/sr (see http://ssc.spitzer.caltech.edu/mips/sat.html). This is the *hard* saturation level at which even the difference image will be saturated. The approximate *soft* saturation limits are given in Table 8.15.

Table 8.15: MIPS 24 μm soft saturation levels

24 μm exposure time	Saturation level in MJy/sr
30 sec	60
10 sec	180
3 sec	600

8.2.2.8 Accuracy, Stability, and Repeatability

As with all other aspects of performance, the stability and repeatability of the 24 μm array is excellent. Photometry is repeatable to 1% or better and the absolute calibration is better than 5% determined (see section 8.1.3.6) using observations of solar analog, A, and K giant stars.

The 70 μm calibration is based (see section 8.1.3.6) on solar analog, A, and K giant stars. Using all the calibration star measurements taken to date, the accuracy of the 70 μm measurements is around 10%. These calibration measurements include stars with predicted fluxes ranging from 30 mJy to 5 Jy on backgrounds between 4 and 25 MJy/sr. One of the calibration stars (HD 163588) is measured every MIPS campaign and the repeatability of this star is better than 7%. The main difference between these two accuracies is likely the result of the nonlinearities still to be corrected (see the MIPS Data Handbook, or the Gordon et al. papers on data reduction, or Rieke et al. 2004, ApJS, 154, 25, all available on the SSC website, for more details).

Unlike the other two arrays, the calibration of the 160 μm data is based on a variety of objects (see section 8.1.3.6) including asteroids. The resulting multiple calibrations agree to within 12%. However, because these were not the planned calibration sources, the 160 μm flux calibration is currently more uncertain than for the other two MIPS arrays. Improvements in the calibration will depend largely on further observations and modelling of asteroids. Repeatability of measurements of the same object is ~10%.

All archived MIPS data will be periodically reprocessed as significant changes in the calibration are determined and implemented. The 160 μm data offer the greatest challenge to the calibration process, and it is likely that those pipelines and data will undergo the most modifications as they continue to improve over the first several years.

Interested readers should also consult sections 8.1.3.5 and 8.1.3.6 on calibration in general and section 8.3.1 on specific calibration procedures, in addition to the MIPS Data Handbook. For further details see Engelbracht et al. (2007), Gordon et al. (2007), and Stansberry et al. (2007; all three papers, PASP, submitted).

8.2.2.9 Observing Time Estimates for MIPS

The total amount of time required to obtain an observation with MIPS will generally be significantly longer than the on-source integration time estimated from the source brightness, background brightness, and the sensitivity of MIPS (see sections 8.2.2.2–8.2.2.5). The amount of overhead for any particular observation is a function of the integration time and the observing mode selected.

Spot keeps track of all of the observing time overhead, using the most recent models available. We provide the following 2 subsections to provide the interested observer with an overview of the factors that contribute to the

observing overheads for the purposes of planning their programs in general terms (the first subsection, 8.2.2.9.1) and to allow for back-of-the-envelope estimations of exposure time (the next subsection, 8.2.2.9.2).

Note that the observing overheads that appear here are approximate, and that the values that appear here are not to be used to supplant the values returned by Spot. Spot provides the highest fidelity estimates achievable on the ground.

8.2.2.9.1   Sources of Observing Time Overhead

Observing time overheads are incurred as a result of both spacecraft and MIPS operations constraints. Spacecraft operations that result in overheads are primarily related to the time required for offsetting, slewing, and settling/pointing acquisition. As a result, programs that minimize the number of pointings will be more efficient than those with many pointings. Some programs will be able to get around this limitation by using scan map mode, which is designed to survey large areas without incurring excessive overheads associated with slew and settle. However, scan map mode has integration-time limitations, and may not be suitable for particularly faint sources. In photometry/super resolution mode, telescope offsets are required in order to switch between the three MIPS bands; each such offset and settle requires roughly 40 seconds to accomplish. In addition to slew and settle times specifically required within a particular program, some amount of slew time must be included for switching between observing programs. A telescope slew overhead of 215 minutes is also added to each AOR (an additional 300 sec overhead is charged for moving targets).

There are two primary types of observing time overhead resulting from MIPS operations. The first of these has to do with the operation of the on-board calibration sources, the stimulators (or 'stims'). The stims are employed very frequently (several times per observation cycle) for both the 70 μm and 160 μm arrays in order to track the detector response carefully (see, e.g., sections 8.1.3.5.2, 8.2.1.1, and 8.2.1.2). This is true for all of the MIPS observation modes. Observers have no direct control over the frequency of stim flashes. Observations made in the 24 μm band are not currently subject to a stim flash overhead. The second type of overhead associated with MIPS operations involves dither motions. In photometry mode, all observations include dithers to place the source of interest (or the desired pointing position) at a number of places on the array. Most of the dither motion is accomplished through the use of the scan mirror mechanism, which moves virtually instantaneously relative to MIPS integration and readout times. However, the dithers also involve at least one nod of the telescope in the direction perpendicular to the scan mirror motion direction, and these nods carry an overhead.

The photometry mode dithers cannot be circumvented by observers. These required dither patterns are designed to improve the accuracy of MIPS photometry by placing sources at several positions on the detectors (both 70 and 160 μm arrays), and to provide a filled 2-D image for the 160 μm array. While

the dithering does not result in an observing overhead per se, it does mean that the minimum integration time on a source is a function of the product of the observer-selected exposure time and the number of dither positions. For bright sources, these minimum integration times might exceed the integration time required by a significant amount. Further discussion of this can be found in the descriptions of the various observing modes in section 8.2.1.

8.2.2.9.2   Fractional Observing Time Overhead

Here we provide illustrative values that can be used for estimating the observing time required to complete an observation given the integration time required to achieve the desired signal-to-noise on a source. The observing times derived here are approximate and subject to small changes. *Observers are required to use time estimates provided by Spot when planning and submitting proposals.*

The amount of overhead incurred in making an observation can conveniently (if approximately) be represented as a factor which, when multiplied by the desired integration time, gives the observing overhead time. The total observation time (or total observation "wall clock time") is just the integration time plus the overhead time, and can be represented as:

Observing Time = (1 + overhead factor) × Integration Time.   (8.3)

Table 8.16: Approximate overhead factors for MIPS photometry/super-resolution (small field option)

Exposure Time (MIPS sec), Band	First-Cycle Integration Time (sec)	First-Cycle Overhead Factor	Seven-Cycle Integration Time (sec)	Seven-Cycle Overhead Factor
3 sec, 24 μm	48	1.40	312	0.36
10 sec, 24 μm	166	0.41	1046	0.11
30 sec, 24 μm	501	0.13	3144	0.04
3 sec, 70 μm (coarse scale)	38	2.07	226	0.68
10 sec, 70 μm (coarse scale)	126	0.85	881	0.38
3 sec, 70 μm (fine scale)	25	4.47	176	2.03
10 sec, 70 μm (fine scale)	84	2.57	587	1.53
3 sec, 160 μm	6	22.13	44	12.19
10 sec, 160 μm	21	12.99	147	8.39
3 sec, 160 μm (enhanced)	13	30.6	107***	11.0***
10 sec, 160 μm (enhanced)	45	14.1	357***	8.9***

***Cycles 7, 11, 13, 17, and 19 are not allowed in 160 μm enhanced mode. The numbers given in the table are for 8 cycles.

In Table 8.16, we tabulate the overhead factors for observations taken in photometry/super resolution mode for the available MIPS exposure times. Note that the tabulated integration times are in real time seconds, not "MIPS seconds" (see section 8.1.4.1). While the combination of dither patterns and minimum exposure time results in fairly long minimum effective integration times, almost any arbitrary integration time longer than the minimum can be obtained by combining sets of 7-cycle observations with single-cycle observations. For most dither patterns, for cycles up to 7 there is only one spacecraft offset, with all other dithers being done with the scan mirror. This is to minimize accumulated overhead caused by unnecessarily repeating motions that require slew and settle overheads, while scan mirror motions do not. For more than 7 cycles, the number of spacecraft motions depends on the natural break points in the sequence. For the purposes of estimation of observing times, the simple method described here will give adequate results. Spot accurately accounts overheads associated with multiple cycles of the observing sequence. Note that for purposes of determining whether a source will saturate the MIPS detectors, the exposure time matters, while for purposes of calculating signal-to-noise ratios, integration time is the appropriate quantity to consider. Note also that these overhead factors are for observations in a single band; if multiple wavelengths are desired, an additional 40 second overhead is incurred for offsetting the spacecraft to bring the object into view of each additional band.

Observing overheads for scan map mode are somewhat more difficult to summarize than are those for photometry mode. Because the three MIPS arrays view different areas of the sky, the length of a scan leg must be somewhat longer than the length of the region being mapped. The overscan distance required to obtain a 3-wavelength full-coverage map is about 20.7′. Data are obtained in the overscan region, but are not co-spatial in all 3 bands, and might have somewhat larger than usual pointing uncertainties in the first several frames. While viable science data, the time spent covering the overscan region is considered observing overhead. The amount of overhead depends on the scan rate selected by the observer, as summarized in Table 8.6.

Exposure times in scan-map mode are determined by the scan rate selected. Likewise, the integration time on any particular point in the scan-map is also determined by the selected scan rate, although deeper integrations can be obtained by re-scanning a region and coadding the resulting maps. The exposure and integration times for the three MIPS arrays as a function of scan rate are detailed in section 8.2.1.1, and that section should be referred to in order to estimate sensitivity as a function of flux, and saturation fluxes as a function of scan rate. The length of time required to complete a 5′ wide scan leg is just the length of the leg divided by the scan rate (including the overscan distance). If the region to be mapped is wider than 5′ (2.5′ for full 70 μm coverage), the map must be built up by specifying multiple scan legs, including specifically some overlap between those legs to ensure complete coverage of the region at all wavelengths. Each time a scan leg finishes and a new scan leg is begun, an overhead of about 30 to

60 seconds for "turnaround time" (time to stop the scan and begin set up of the new scan leg), which includes an offset and settle, is incurred. There is also a 30 second overhead incurred at the beginning of each scan map AOR for establishing initial pointing, settling, and beginning to slew.

8.2.2.10   Image Quality and Distortion

For most applications, MIPS can be taken to be purely diffraction-limited. The predicted wavefront error introduced by the instrument optics at the center wavelength of each photometric band is listed along with the requirements in Table 8.17. Only in the 24 µm band is the predicted wavefront distortion significant, and the quoted value holds only near corners of the field of view; the wavefront error is significantly smaller over most of the array. The distortion is measured to be exactly as predicted at 24 µm. Although the distortion at 70 and 160 µm is yet to be measured in detail, early results suggest that the values are not inconsistent with the predictions.

Because the MIPS optical train is made up purely of off-axis reflective elements, some degree of scale change across the re-imaged focal planes is inevitable. All optical trains meet the requirement that this distortion should not exceed 10% (defined as $100 \times [(\text{max scale})/ (\text{min scale}) - 1]$). Most bands meet the requirement comfortably (see Table 8.17). Nonetheless, to coadd images taken at differing places on the array, and for other science applications, it is important to correct the data in all bands for the image scale changes.

Nominal corrections of optical distortions are performed in pipeline processing of the MIPS images and maps, and necessary data can be supplied to the observer to perform the corrections on the Basic Calibrated Data (BCD) images if desired (see Data Products, section 8.3.2).

Table 8.17: Optical distortion in MIPS

	RMS wavefront req./performance	Distortion (req. 10%)
24 µm band	0.067/0.061	4.7%
160 µm band	0.040/0.021	9.7%
70 µm survey	0.040/0.022	1.9%
70 µm super resolution	0.040/0.021	7.5%
70 µm SED	0.040/0.022	4.1%

8.2.2.11   Scan Alignment

Scan Map Mode operation requires the use of the MIPS scan mirror to compensate for image motion during Spitzer spacecraft continuous scan motion. The desired result is the simultaneous delivery of a succession of stationary sky images to all MIPS FPAs.

Ideally the spacecraft slews its view vector along a straight line in object space

(able to be oriented to MIPS needs) at a constant angular velocity, while the MIPS scan mirror moves at a constant angular velocity (to which the spacecraft rate can be matched). Given an ideal optical system, the full-field images will remain fixed on the FPAs for the duration of the integrations.

The actual instrument departs slightly from these ideal conditions. There are two general classes of image smear due to the scan mirror motions:

A. Image smear sources in individual bands
   1. Departures from constant angular velocity.
   2. Changes in field distortion or magnification during scan motion.
B. Image smear sources as a result of a mismatch between bands
   1. Angular magnification mismatch between the sides of the instrument at the scan mirror.
   2. Direction of image motion mismatch between sides of the instrument.

Analysis of scan mirror operation during instrument tests indicates that all of these sources of image smear should lie well within the FWHM of the diffraction limited image, and for most observing modes should be virtually undetectable in the observed shape of the PSF. The one minor exception is in the scan map mode at the fast scan rate where a slight elongation of the images may occur due to sources of type B, which cannot be reconciled through adjusting the spacecraft motion or the scan mirror action. The mismatch may become apparent in this mode because of the relatively long throw over which the scan mirror is used to freeze the image motion.

8.2.2.12   Ghost Images

Reflections from filters in front of focal plane arrays frequently cause ghost images. In the case of the 160 μm band, it has been possible to tilt the bandpass filter far enough so that no ghosts will occur, but at 24 and 70 μm very faint, slightly out-of-focus ghost images do appear. The intensity of the ghost images measured during tests of the integrated MIPS is $\leq 0.5\%$ of the brightness of the primary image at 24 μm, and is 2% of the brightness of the primary image at 70 μm. The single ghost image formed on the 70 μm array is 6 pixels (in wide-FOV mode) from the primary image.

For additional information, please see the MIPS Data Handbook.

8.2.3 MIPS AOT Examples

In this section, we present information relevant to filling out Astronomical Observing Templates (AOTs) for basic MIPS observations. The MIPS AOT inputs are relatively few, and designing a typical MIPS observation is not complicated. However, the design of all MIPS observations should be based on a careful examination of the far-infrared properties of the particular region of the sky to be observed, and a good understanding of MIPS operation, capabilities, and limitations. In addition, the orientation and rotation of the Spitzer focal plane can

be a factor, as can the three-hour limit of a MIPS AOR (e.g., for large scan maps).

There are four major operational modes (or Astronomical Observing Templates; AOTs) of the MIPS instrument: (1) Scan Map (includes slow, medium and fast scan rates, where incomplete coverage at 160 μm is obtained at the fast scan rate); (2) Photometry and Super Resolution (includes super-resolution, large source, and small source options, and can be used to obtain multiple images through the cluster and raster-map options); (3) Spectral Energy Distribution (SED; can be used for a single spectrum or step and stare spectral mapping); and (4) Total Power Mode (used to obtain absolute brightness measurements for highly extended sources).

Observing in any of these modes involves the acquisition of multiple data frames, not just a single frame. The multiplicity of frames provides for rejection of cosmic rays, calibration of the Ge:Ga focal plane array data in the light of the two time-constant behavior of those detectors, adequate sampling of the point spread function (PSF; especially for super resolution observations), and, in the case of the 160 μm array, for building up a filled image using multiple offset exposures of that 2×20 pixel array. The total number of images obtained depends on the total integration time needed, the observational mode, and the array. The multiplicity of frames is prescribed within the AOTs, and the observer selects the number of times to repeat the basic pattern of a template in order to build up the desired integration time. The exposure times for each frame are also limited by the AOTs, with the result that the total integration times that can be specified are quantized.

The Cluster Option is a way to conveniently specify repetitions of one or multiple observations at several closely grouped pointings (within approximately 1° of initial target). The observations done at each cluster position are identical, and done in the same pointing sequences as for a single target AOR. The observational modes that can be specified under the cluster option are MIPS Photometry/Super Resolution, SED, and TP.

Observations of moving targets can be made in a co-moving frame for all MIPS AOTs, including Scan Map.

For all AOTs, the *number of cycles* selection is the number of times to execute the basic observing sequence or full map sequence as described in section 8.2.1. Observers are reminded that in photometry/super-resolution mode, MIPS always takes multiple, dithered frames, each with the specified exposure time. Because of this, the total exposure time that can be specified by the user is an integer multiple of the single frame exposure time, and the number of dither positions imposed on all photometry AORs; the total on-source integration time is quantized as a result. The number of cycles selection determines how many of these integration-time quanta will be obtained. Likewise, on-source integration time is quantized in scan-map mode, depending on the selected scan rate, and selecting more than one scan cycle will build up integration time in multiples of

that single-cycle integration time.

Observers can specify that a small region (a raster map) be covered in photometry mode. The entire AOT cycle is repeated at each raster map location; the number of *map cycles* is different and refers to the number of times the whole map is repeated. The exposure time per pixel reported by Spot refers only to the times within one AOT cycle; observers should use depth-of-coverage maps (see the Spot User's Guide) to further determine how many visits to their target are obtained in the raster map. (See for example the 160 μm raster map example as part of section 8.2.3.3.)

Because of noise improvements relating to characteristics of the stressed Ge:Ga 160 μm detectors, when a 10 sec exposure time is selected, the 160 μm detectors will be reset twice within the exposure interval. This effectively results in two exposures at 160 μm before dithering when using the 10 sec exposure option. The 160 μm sensitivities for 10 sec exposures include the multiple resets, so no impact to observation planning arises from the implementation details.

8.2.3.1 Best Observing Practices: MIPS in General

In addition to the items mentioned above, observers should also take into account the following items. Latent images at 24 μm created by bright sources could affect data, and are not currently removed by the SSC pipelines (see section 8.1.3.3.2). There is a short-wavelength light leak at 160 μm; extragalactic objects are unlikely to be affected, but stellar programs might be, especially debris disk programs looking for subtle excesses (see section 8.1.3.6.6). A MIPS second is slightly longer than a true second (see section 8.1.4.1); plots and tables in this chapter and on the MIPS SSC website are in MIPS seconds unless otherwise specified, but times returned by Spot are true seconds. For far infrared Ge detectors, observers should make sure they take enough background (source-free) data. For large sources we recommend using the chopping large-field photometry mode or scan mapping. At 160 μm we recommend using the new 160 μm enhanced photometry mode for small sources. At least four images of a 160 μm target are recommended to ensure reliable photometry. Images of objects near the ecliptic plane are likely to include asteroids and/or Kuiper Belt objects; a second observation at a later time is necessary to ensure that moving targets have been removed. Observations at MIPS wavelengths may be limited by noise (including confusion noise; see section 8.2.2.1), but observers should ensure sufficient data redundancy for cosmic ray rejection, even if the confusion limit is reached with shorter exposures. It is best to build up long integrations with long exposure times, e.g. large total integration times built from 3 s exposures will not produce as high-quality a final product as fewer cycles of 10 s exposures.

Observers should *always* visualize their observations using the latest version of Spot to be sure that the observation is doing what was intended.

Also see additional "best observing practices" customized for photometry (section 8.2.3.5), for scan maps (section 8.2.3.7), and for SED (section 8.2.3.9).

8.2.3.2 Example: Photometry AOT

Assume a photometry observation in the three MIPS bands is desired to determine the far-infrared colors of a point source, or one that is only slightly extended ($\leq 2'$ at 24 and 70 µm, $\leq 1.5'$ at 160 µm), such as a circumstellar disk source. From the observer-predicted flux estimate for the three bands and the MIPS sensitivity charts, it is determined that 200 seconds of on-source integration time is needed at 24 µm, and 100 seconds is required at 70 and 160 µm. This will require two cycles of the basic dither pattern at 24 µm (278 seconds total on-source integration time for the 10 second exposure time), one dither-pattern cycle at 70 µm at the default pixel scale (120 seconds on-source at the 10 second exposure time), and five dither-pattern cycles at 160 µm (giving 100 seconds integration for the 10 second exposure time). Figure 8.36 shows an example of the AOT front-end for specifying this observation in the MIPS Photometry/Super Resolution section of Spot.

Observers are reminded that two cycles of the 160 µm small field photometry AOT provide only a minimum level of data redundancy (4 complete images). Spot will alert you if you have fewer than the recommended number of cycles.

For very deep integrations, observers may encounter the photometry/super resolution cycle number limit (20 cycles) before the maximum AOR time (3 hrs) is reached. For such observations, it is recommended that the observer reposition the object or field on the array and repeat cycles until the desired integration time is reached. This can be most easily accomplished by creating cluster offset targets or by doing small raster maps; for example, if a total of 30 cycles is desired, create a cluster target with two positions close together (or do a small raster map) and ask for 15 cycles at each cluster position.

Figure 8.36: Spot AOT front end showing the inputs required to specify the basic photometry observation in three colors of a point-like source.

8.2.3.3 Example: Super Resolution AOT

To obtain data that can be post-processed for super resolution requires a somewhat different observing strategy from normal photometry observations. The 24 μm pixel sampling provides proper sampling of the PSF using the standard photometry dither sequence, so no modifications are necessary. For super-resolution at 70 μm, the fine pixel scale must be selected, while at 160 μm, a 3×1 raster map should be selected, (see Figure 8.38) using 1/8[th] array stepping in the cross-scan direction, to provide the needed pixel sampling.

If, for example, the observer estimates that 600 sec of on-source integration time in 10 sec exposures is required at 70 μm in the fine scale, then at least eight cycles of the standard fine scale observing sequence will be required. If, as above, 100 seconds is required at 160 μm, three cycles of the standard sequence at 10 second exposure times will provide about 180 seconds per pixel on-source integration time in the 3×1 map sequence. Note that by selecting "Calc. Obs. Time" in the AOT window, Spot shows the integration time *per pixel* for the 160 μm array as about 60 seconds; at this point, Spot does not know where the source is on the array. Since the source is moved 3 times on the array by only 1/8[th] array steps each time in this example, the total integration time per pixel *on source* is about 180 sec.

In Figure 8.37 and Figure 8.38, example Spot inputs are shown that will obtain super resolution sampling in the three MIPS bands.

Figure 8.37: Spot AOT front end showing the settings for defining a super-resolution observation at 24 μm and 70 μm. In order to obtain the proper sampling at 160 μm, the "Set Map" checkbox for 160 μm must be activated.
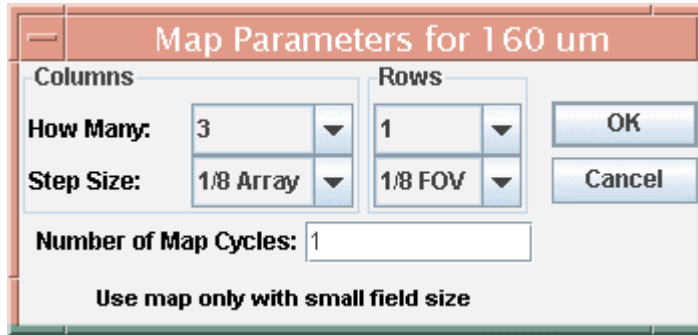
Figure 8.38: The 160 μm raster map definition dialogue box, showing the settings required to obtain a super-resolution observation by setting up a 3 by 1 raster map. Note that "columns" defines the number of FOVs in the cross-scan direction, and "rows" defines the number of FOVs in the scan direction.

8.2.3.4 Example: Large Field Photometry and Raster Mapping

Suppose an observer wishes to image a face-on spiral galaxy that is about 2′ in diameter, and there is fairly uniform sky around it. The standard photometry dither patterns are designed to image sources less than about 1′ in size. The "large field" patterns are suitable for sources less than about 2′ in size. The large field dither patterns (described in section 8.2.1) assemble images centered on the source coordinates and acquire background data near the source in the "in-scan" direction. For 24 μm, the background data are acquired with a telescope slew to an observer-defined distance from the source. At 70 and 160 μm, the sky data are acquired through scan mirror motions. Specifying a large field source is no different from standard point source photometry, except that a background chop distance for 24 μm needs to be specified (300″ is the default), and the large field option selected. See section 8.2.1 for how the data are acquired. Figure 8.39 shows an example large field specification for a galaxy like that in Figure 8.27.
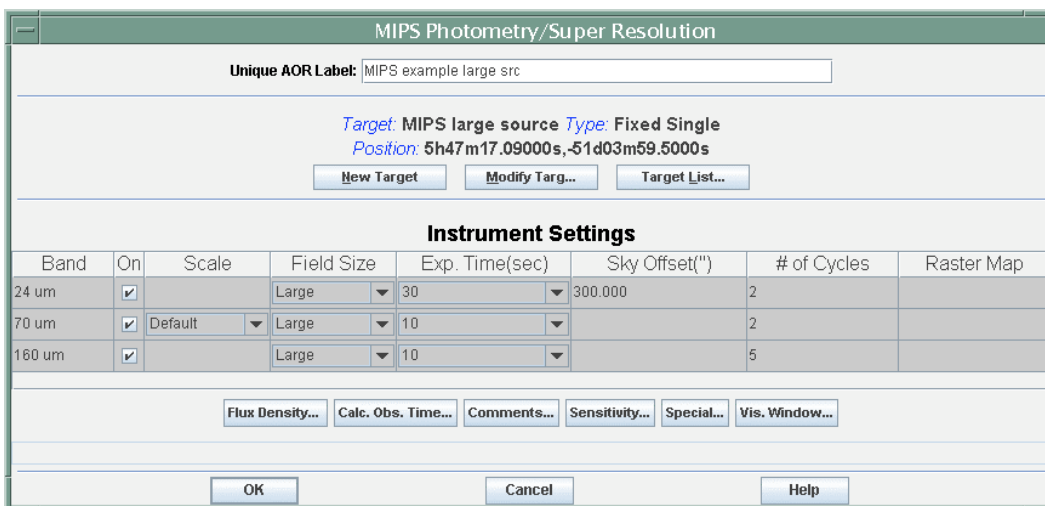

Figure 8.39: Spot large-source AOT dialogue box showing the settings that would be selected for imaging the ~2′ diameter galaxy of this example.
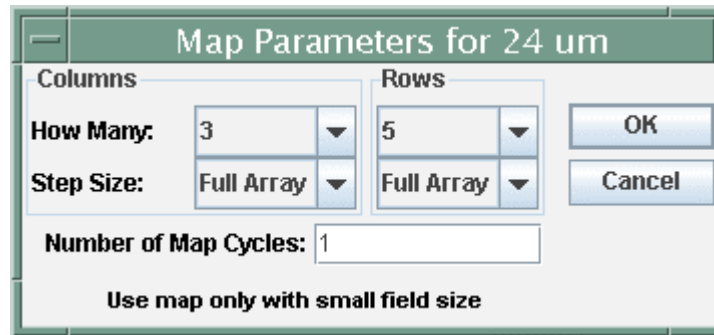
Figure 8.40: Sample input window for a simple raster map at 24 µm. Note that "columns" defines the number of FOVs in the cross-scan direction, and "rows" defines the number of FOVs in the scan direction.

For regions with complex background, or sources larger than about 2′, the photometry raster map option can be selected. This option is only available with the small field dither patterns and the default 70 µm pixel scale (10″). Stepping is done in array coordinates only, in the cross-scan direction and the in-scan direction. Steps in the in-scan direction at 160 µm are defined in fractions of that detector's effective field of view (2′×5′), which is filled through multiple steps of the scan mirror. All other step sizes are defined as fractional array size steps. Because the small field dither patterns move the pointed position over a relatively large area of the arrays, a MIPS raster map will have large overlaps of exposures. This overlap is required for proper frame-to-frame calibration of these far-infrared Ge:Ga images, especially when complex background structures are present in the images. Examples of the inputs for a simple raster map are shown in Figure 8.40.

There is no simple formula for determining whether a large region of sky can be observed more efficiently using a raster map or a scan map. Scanning is very efficient for covering large regions on the sky in all three MIPS bands, but does not integrate especially deeply. For deep observations of regions 2′ to 10′ in size, it will typically be more efficient to use the raster option. The observer should compare the times required by example inputs to Spot in order to determine the most efficient observing mode. Other considerations should also be taken into account, such as data quality, depth, and full coverage at 160 µm.

8.2.3.5 Best Observing Practices: MIPS Photometry

In addition to the items mentioned immediately above and for MIPS in general (section 8.2.3.1), observers should also consider the following items relevant specifically to MIPS photometry. Super-resolution data need to be post-processed to take advantage of the sub-pixel sampling inherent in the AOT; the SSC pipelines do not currently provide this service.

The AOTs each provide multiple images of the target; observers select the number of times that the 'standard cycle' is repeated. When defining raster maps,

observers select an additional "map cycle" which indicates how many times the entire map is carried out; an AOT cycle is completed at each map location.

In photometry observations, the dither pattern is such that the region of highest-quality S/N and best coverage is less than a full FOV; use Spot visualizations to determine the region of sky actually covered by all pointings for any given AOR.

Because of the smaller effective size of the 70 μm array compared to previous expectations, to obtain the same coverage as pre-launch expectations for large-field photometry, some users will need to efficiently map the approximately 5′×5′ (2.5′×2.5′ fine scale) area covered by the original AOT. In order to do this, we suggest that the user create a cluster target with offsets of (0, +80″), (0, -80″) [for fine scale, (0, +60″), (0, -60″)] in array coordinates observing "offsets only." This will allow recovery of the full area of the original AOT while saving the slew tax from mapping with 2 AORs.

8.2.3.6 Example: 4 Scan Map AOTS

MIPS scan mapping is useful for efficiently imaging large areas of sky in all three MIPS bands. As discussed below, the design of a MIPS scan map will be influenced by several factors; the larger the scan map, the more stringent the constraints resulting from these factors will be. The primary factors to consider are: 1) required integration time (i.e., scan rate), 2) scan leg size, and 3) scan leg orientation. The basic inputs for a MIPS scan map (scan rate, scan leg length, and scan leg offset, each selected from a list in the Scan Map AOT front-end), directly address two of these. The scan leg offsets in the cross-scan direction determine the amount of overlap between scan legs, and can be different for a change from the initial, forward scan direction, than for a change from the reverse to the forward direction.

Definitions: A "scan leg" is defined to be in one direction only. An "up-and-back" scan is made of two scan legs. The scan leg length as input to Spot is the "full coverage" length in all three MIPS bands. The actual distance over which a scan is made is approximately 21′ longer than the full coverage scan length. This "overscan" area does not provide co-spatial coverage in all three bands, but otherwise overscan data will be of the same quality as in the rest of the scan map. The target coordinates are defined to be the map center location. The cluster option is not available for MIPS scan mapping.

The orientation of the scan legs on the sky needs to be taken into account in designing the map. Obviously, orientation will be a consideration for covering the desired map area. However, orientation can become a complicated design issue for large maps because of limitations on Spitzer pointing and roll angle (see section 3.2.5), and because there is a three-hour limit on the duration of a single MIPS AOR (longer observations must be submitted as a series of AORs). The orientation of the Spitzer focal plane is a function of ecliptic latitude and time of year. If a scan map is broken up, the observer must take steps to insure that the

timing of execution of the component scan maps is such that adequate overlap between those component maps is achieved. This is not a significant problem for fields lying near the ecliptic plane, as the focal plane does not rotate very much in that region of the sky. For fields at high ecliptic latitude, the focal plane orientation changes appreciably over a period of a few days, so large scan maps in that region must be carefully designed so that there are no unwanted wedges of unobserved sky between the component maps. On the other hand, maps at high latitude can be specified with an arbitrary orientation by specifying the time at which the observations are to be obtained. This flexibility is not available for low-latitude maps, where map legs all run nearly north-south.

Multiple AORs requiring similar orientation can be chained, or constrained to be done close in time. The observer is advised that specifying very tight constraints will limit the schedulability of the AOR. A large map constructed out of small, square submaps that can be pieced together with acceptable overlap and no constraints will be the most straight-forward to construct and suffer the least chance of scheduling difficulties. (See section 5.5.3 for more information on constraints.)

"160 micron required" check box: Because of Spitzer's need for active cooling of the telescope primary mirror and single instrument operations, there will be (known) times when the telescope is warmer than ~6 K, thereby limiting 160 μm performance. Under normal operations, the telescope will be maintained at ~6 K during those MIPS campaigns containing observations requesting 160 μm, and high-quality 160 μm data will be obtained. For observations that happen to be scheduled immediately after an instrument changeover, the highest level of performance cannot be assured without allowing extra time for cooldown. In order to maximize Spitzer's cryogenic lifetime, starting in late 2005, MIPS campaigns will be organized into "warm" (no 160 μm observations possible) or "cold" (160 μm observations possible) campaigns. Programs using scan maps and requiring MIPS 160 μm observations must justify this specifically in the proposal. Proposers should carefully consider the science objectives of their program before requesting 160 μm data. For observations requiring specialized scheduling or strong constraints, this might negatively impact scheduling of the observation (or make it entirely impossible to schedule). If the box is not checked, high-quality 160 μm data should be considered serendipitous.

Fast reset option: As described earlier (section 8.1.4.3), the 160 μm array is normally reset twice per integration for 10 second integrations. Possible integration times in the Scan Map AOT are 3, 4, and 10 seconds; the 160 μm array is automatically reset at 5 second intervals for 10 second integrations and at 2 second intervals for 4 second integrations. There will be observations for which the background at 160 μm could saturate the detectors even for 3 second integrations (the fast scan rate), while the 70 μm background will not limit the observation in this way. By clicking the Fast Reset button in the Scan AOT, the observer causes the 160 μm array to be reset every 1.5 seconds in fast scan only

(there is no effect at medium and slow scan). This option should only be used when required for high background regions.

Backward mapping: Some observing scenarios, such as 70 or 160 μm studies of large extended galaxies, favor mapping in one direction and then remapping again in the reverse direction. This option exists for just such occasions.

Figure 8.41: Spot AOT front end for 1°×0.5° medium scan rate map. This AOR design only provides single sampling at each point at 160 μm.

Examples of simple maps: **Figure 8.41** shows the Spot configuration for a basic medium scan rate, 1°×0.5° map, with 160″ (half-array) cross-scan steps. The total

time required to obtain this map is very nearly 3 hours, so only one map cycle can be done per AOR. For this configuration, the 160 µm scans will have minimal overlap and redundancy. It is therefore recommended, for good quality 160 µm data, that this AOR be repeated at least once. Inherent multiple redundancy at 24 and 70 µm removes that need if only these bands are of interest. However, detection of asteroids in these data will require at least a second map obtained at a later time.

Figure 8.42: Spot AOT front end for the 1°×0.5° fast scan rate map with nearly full coverage at 160 µm provided by the ¼-array cross-scan step for return legs.

Figure 8.42 shows the Spot inputs to define a 1°×0.5° map using the fast rate. The

160 μm data will contain coverage gaps (see section 8.2.1), but this AOR design nearly fills the gaps because the small cross-scan step for the return legs of the map brings the second leg back along nearly the same path, but offset by a ¼ array width so that the missing readout does not hit the same patch of sky more than once. A full-pixel-sized offset is made automatically to achieve a fully sampled scan path. The minimum scan leg length for the fast rate is 1°.

Figure 8.43: Spot AOT front end for the 5°×5′ scan map which is rescanned 8 times to give good coverage at 160 μm.

Example of a narrow, multi-scan strip: A scan map might provide a very effective and efficient way of acquiring multi-band data on a long, narrow strip of sky. An

example might be repeatedly scanning in a narrow strip to build up the desired integration time. An example of a scan strip 5º long rescanned 8 times is shown in Figure 8.43. This particular map design will fully sample the 160 μm field four times. A quarter-array-width cross-scan offset is made for enhanced sampling and protection against bad pixels.

8.2.3.7 Best Observing Practices: Scan

In addition to the items mentioned immediately above and for MIPS in general (section 8.2.3.1), observers should also consider the following items relevant specifically to MIPS scans. Step by at most ½ array steps to obtain full sky coverage at all 3 wavelengths. For the best data at 70 μm we recommend cross-scan steps of ≤ 148". For highest S/N coverage, step by at most ¼ array steps. The AOT provides inherent redundancy in all wavelengths, but there are more images of a given source at 24 and 70 μm than at 160 μm. Fast scans provide only half coverage at 160 μm, so the scans must be repeated to approximately fill in the gaps. To achieve the recommended minimum 4 images of the same patch of sky, observers should plan additional scans with small cross-scan steps.

Bright objects may leave latents in the data; to limit these effects, try to arrange for covering the same region of the sky in two directions (e.g., forward and back). Extremely bright objects might be isolated in a separate AOR.

Figure 8.44: SED AOT front end for SED example.

Note that the first several frames of a scan leg necessarily have extrapolated stim backgrounds (see Data Handbook for more information on this effect). In general, but especially in the case of fast scans over bright objects, be sure to have enough background measurements around your source that extrapolated stims are not an issue.

8.2.3.8 Example: Spectral Energy Distribution

The observing sequence for MIPS SED is simple and straightforward. Consider a relatively small target (<1′) source. In this case, a chop distance of 1′ will reach sufficiently far off-source to sample the sky. In Figure 8.44, we have selected this chop distance. A ten-second exposure, repeated in 4 cycles, yields on-source integration time of about 250 seconds. Note that in a single cycle, spectra are taken at two positions in the slit for a total of 6 on/off pairs of observations per cycle; 4 cycles repeats this sequence 4 times. A simple raster map option similar to that for Photometry/Super Resolution (see above) is available as well.

Figure 8.45: TP AOT front end for TP example.

8.2.3.9 Best Observing Practices: SED

In addition to the items mentioned above for MIPS in general (section 8.2.3.1), observers using MIPS SED mode need to consider their exposure times and likelihood of saturation. Since the saturation limits are so high for the 70 μm array, it's rare to saturate SED mode. Thus, except in situations where extremely bright objects are being observed, one should always use the 10 second exposure

time.

8.2.3.10   Example: Total Power Mode

The observing sequence for MIPS TP is also simple and straightforward. Figure 8.45 shows the Spot front end for a TP observation. In this case, we have selected 1 cycle of 3 s integration at 24 μm, and 10 s at both 70 and 160 μm.

8.3 Data

8.3.1 Data Calibration

8.3.1.1 Wavelength Calibration

The SED wavelength calibration was obtained on-orbit by observing the planetary nebulae NGC 6543, NGC 6826, and IC 3568. In addition, the supernova remnant IC 443 and the H II region 30 Dor were observed. For all objects, except IC 443, the [O III] 88.3 μm emission line was detected and measured. The [O III] line was the only spectral feature identified in these spectra. For IC 443, the only identified spectral feature was [O I] 63.2 μm. Combining the measurements from these nebulae yields a spectral coverage of 52 μm to 97 μm for SED mode and a dispersion of 1.71 μm/px. There is a shift of less than a pixel in the position of emission lines along the usable length of the slit, and the centroid of the spectrum on the array shifts by about a pixel from the blue to the red end of the spectrum. These slight distortions will be corrected in the pipeline data reduction.

Ultimately, it is believed that the wavelength calibration will be accurate to better than 3% in resolution and absolute wavelength. While ongoing adjustments of the calibration are possible, SED data frames will be rectified such that all columns in a FITS frame have the same wavelength, and header fields for calculating the wavelengths will be supplied with the data.

8.3.1.2 Science Observation Calibration Procedures

As described in section 8.1.3.5, the internal calibration sources, or stimulators, play an all-important role in calibrating the MIPS data. The stimulators are used as relative calibrators, tracking drifts in the responsivity of the detectors. The brightness of the stimulators is tied to periodic observations of well-calibrated celestial standards. At a minimum, the celestial standards are observed at the start and end of each MIPS campaign using the standard Photometry AOT (see next section).

The fundamental MIPS flux calibration is against normal stars (section 8.1.3.5). Except for the occasional dedicated calibration AOR of celestial standards, an individual science AOR does not depend on any other science AORs for obtaining calibration. Each MIPS AOR is internally calibrated. Because of the use of the stimulators as relative calibration sources, this method is robust and stable over an instrument campaign. This relative calibration method also allows the varying instrument response to be frequently referred to the signals from the celestial

standards. For the Ge:Ga arrays, the stimulators are flashed approximately every 2 minutes or less as an integral part of the basic observation sequences. Because of the inherent longer term stability of the Si:As array, that detector uses much less frequent stimulator flashes. The frequency of stimulator flashes is determined by the SSC and implemented in the AOT design, and cannot be selected by the observer.

Detector dark current measurements are made for each array individually by positioning the scan mirror such that the array views only the interior of the cold instrument. The optical layout of MIPS is such that only one array at a time can be completely hidden from light entering through the telescope. In fact, only about two thirds of the 70 μm array can be placed in the dark at once, so for that array the dark current measurement is made in two pieces. Dark current measurements are expected to only be required at the beginning and end of each MIPS campaign. To avoid any scattered light from nearby very bright sources affecting the measurement, the telescope is pointed toward a region known to be free of such sources during the dark current measurement sequence.

Flat-fielding of the Ge:Ga detector arrays is achieved by using calibrated stimulator illumination patterns. The stimulators are calibrated using a flat field matrix determined from "uniform" sky measurements at the start and end of instrument campaigns. The flat field procedure takes advantage of fundamental MIPS observation techniques to obtain the highest quality flat field data possible (median combine of multiple scan map legs with source rejection, for example). Standard astronomical observation practices are used in determining the flat field response of the arrays and in the processing of the data.

8.3.1.3 Observing Campaign Calibration Procedures

Sets of basic data for determining the absolute calibration of the detectors are obtained at the beginning and end of each MIPS instrument campaign. Various types of well-studied astronomical sources are observed for this purpose (see section 8.1.3.5). It is expected that these observations do not need to be taken more frequently because of the stability of the stimulator output over time, and the exhaustive use of stimulator flashes built-in to all MIPS observing templates. Changes to the specific details of these Start-up and Shut-down sequences can be made at any time without notice. To keep the observatory observation efficiency at a maximum, "unnecessary" calibration data acquisition will be avoided.

Because there is contamination on the pick-off mirror that feeds MIPS (see the Data Handbook for more information), each campaign includes sets of scan-mirror-position-dependent flat fields for the 24 μm array. When we have built up a complete library of all of the possible contamination spot positions, we will not need to devote time to acquiring these flats in each campaign.

8.3.1.3.1   Start-up Sequence
- If not already on for IRS, switch on the Combined Electronics and other

MIPS electronics requiring specific turn-on commanding.
- After the required warm-up time (0.5 hour), proceed with routine calibration measurements.
- Obtain Ge:Ga detector bias calibration observations
- Slew to MIPS "dark sky" region.
- Thermal anneal all detector arrays.
- Acquire dark current data frames for all arrays.
- Slew to MIPS "flat sky" and acquire flat field frames to provide average uniform illumination.
- Acquire flat field projector and stimulator data for the flat field matrix.
- Slew to Observatory cross-calibration flux standard(s) for 24 μm. Observe using the Photometry AOT.
- Slew to MIPS 24 μm flux standard(s) (if different from above) and observe using the Photometry AOT.
- Slew to MIPS 70 μm flux standard(s) and observe using the Photometry AOT.
- Slew to MIPS SED flux standards(s) and observe using the SED AOT; also obtain an illumination correction.
- If not a MIPS-warm campaign, slew to MIPS 160 μm flux standard(s) and observe using the Photometry AOT.
- Proceed with science observations as planned for the instrument campaign.

8.3.1.3.2   Shut-down Sequence
- End science observations at the scheduled time.
- Slew to MIPS "dark sky" region.
- Thermal anneal all detector arrays.
- Acquire dark current data frames for all arrays.
- Slew to Observatory cross calibration flux standard(s) for 24 μm. Observe using the Photometry AOT.
- Slew to MIPS 24 μm flux standard(s) and observe using the Photometry AOT.
- Slew to MIPS 70 μm flux standard(s) and observe using the Photometry AOT.
- Slew to MIPS SED flux standards(s) and observe using the SED AOT; also obtain an illumination correction.
- If not a MIPS-warm campaign, slew to MIPS 160 μm flux standard(s) and observe using the Photometry AOT.
- Obtain Ge:Ga detector bias calibration observations.
- If not required by IRS, switch off the Combined Electronics and other MIPS electronics requiring specific turn-off commanding.

8.3.2 MIPS Data Products

The nature of the current SSC pipeline data processing of MIPS data and the resultant data products are briefly described here and in the following section. More detailed discussion, including examples (even software), can be found as

part of the MIPS Data Handbook. As analysis of data from flight and the flight spare arrays proceeds, revisions to this description and that in the Data Handbook will be made accordingly. This description applies to data obtained and processed within at least the first year of the Spitzer mission, and so applies primarily to Legacy Science and Guaranteed Time observations.

In order to fully understand and assess the characteristics of MIPS data, the data processing capabilities, and limitations for the Ge:Ga data in particular, it is important for the observer to carefully examine and understand the photoconductor detector behavior (section 8.1.3.3), calibration strategy (sections 8.1.3.5 and 8.3.1), data acquisition methods (section 8.2), and how these relate to each other. Gordon et al. (2004, PASP, 117, 503; available on the SSC website) has more details about the data processing algorithms than are included here.

MIPS data are delivered to observers as FITS image files through the Leopard archive interface; observers log into the archive interface to obtain their data. Because of the complexity and redundant nature of the MIPS AORs, three distinct types of data are generated by SSC pipeline processing. Observers will be able to select which kinds of data they would like to download through the archive interface.

Basic Calibrated Data products (BCD) for MIPS are data derived from a single data collection event (a DCE, or a single frame exposure). The Basic Calibrated Data are planned to be the most reliable product achievable by automated processing. BCDs for all AOTs are in the form of FITS images; the SED BCDs are dispersion-corrected and calibrated two-dimensional spectral FITS images. Ge BCD images are corrected for the effects of cosmic ray impacts; cosmic rays are flagged (but not removed) for Si. The total integration time for the individual BCD frames is the observer-specified exposure time (i.e., 3, 4, 10, or 30 seconds), not the total integration time accumulated during the multiple on-source exposures that result from an AOR. No background or sky subtraction is made. BCD products are composed of 2-dimensional FITS image files and a full set of header information keywords, including distortion coefficients, epoch, pointing information, and pipeline and calibration versions. The BCD products are calibrated in MJy/ster.

Post-BCD products (post-BCD) are derived from a full AOR (e.g., a scan map or a dithered photometry observation). These data are the result of combining all individual BCD frames from a single AOR; in the case of photometry mode, the product is an averaged and registered single image suitable for photometric measurements, and in the case of a scan map, the product is a registered mosaic, with first order removal of seams between the component images. MIPS post-BCD products are calibrated in MJy/ster. The post-BCD products are delivered to the observer in the form of a single FITS image file per AOR, and include a full complement of header keywords. **As of this writing the BCD products are, for the most part, what you should start with when working with MIPS data.** For Ge data, in particular, the post-BCD data products are a good way to get an

overview of your data; however, to do science, you should revert to the BCDs and make your own post-BCD products.

The *Raw Data* products are the unprocessed array images (in unprocessed counts per pixel). These FITS image files allow the observer direct access to the data, but still contain the difficult-to-calibrate detector behavior inherent to Ge:Ga detector technology for the 70 µm and 160 µm frames. No cosmic ray removal will have been performed.

8.3.2.1 Calibrated Data Units

The BCD product, which is the primary data returned to observers after pipeline processing, consists of individual frames where the pixel values are in units of MJy/sr. Jansky is a flux density unit defined as:

$$1 \text{ Jansky} = 1 \text{ Jy} = 10^{-26} \text{ W m}^{-2} \text{ Hz}^{-1} = F_\nu \tag{8.4}$$

The conversion between Jansky and flux density in W m^{-2} per unit wavelength is accomplished via

$$F_\nu \times 10^{-26} \times c / \lambda^2 = F_\lambda \tag{8.5}$$

where the wavelength bin-width is specified in the same length units as λ and c. For example, if c is taken as 3×10^{14} µm s^{-1} and λ is specified in µm, the above equation results in F_ν being in units of W m^{-2} µm^{-1}.

8.3.3 MIPS Data Processing

Gordon et al. (2004, PASP, 117, 503) and the MIPS Data Handbook (both available on the SSC website) have more details about the data processing algorithms than are included here.

Basic Calibrated Data (BCD) images are accompanied by additional FITS files. These include, but are not limited to: traceable uncertainties per pixel (an error image), calibration and pipeline reduction pedigree, and data masks for pixels flagged for various data quality criteria.

Production of the Basic Calibrated Data includes removal of electronic signatures (as applicable) such as a "droop" correction in the 24 µm array data (see also Chapter 7), dark current subtraction, non-linearity of sample ramp slopes (Ge:Ga), a robust fitting of ramp slopes to determine total counts per second (Ge:Ga), and the flagging and removal of single or multiple cosmic ray hits using robust Bayesian techniques for the fully sampled Ge:Ga exposure ramps sent to the ground. When possible, corrected ramp segments are adjusted and averaged as required. The inherent redundancy of the MIPS data acquisition adds additional compensation for cosmic ray effects in the combined data. Bad pixels are identified and masked as appropriate. For a uniform diffuse background, the pixels in a BCD image report equal values (within noise statistics; a.k.a. flat fielding of the arrays is performed). The flat field of the Ge:Ga array data is determined and maintained temporally for the Ge:Ga arrays as described in

section 8.1.3.5. Correction of raw counts to flux density per pixel is done based on the detector behavior described elsewhere in this chapter (section 8.1.3.5). The flow, or order, of pipeline processing steps appears in Figure 8.46 (the BCD Si pipeline), Figure 8.47 (the BCD Ge pipeline), and Figure 8.48 (the post-BCD pipeline). These pipelines will continue to be refined during operations.

The Final BCD products for 24 µm data include (but are not limited to) the following items: slope image, difference image (all zeroes except for pixels saturated in the slope image), replaced image (identical to slope image, except non-zero difference pixels are replaced), error images, and mask images. Note that the 24 µm pipeline does not correct for latent images, and although it flags cosmic rays, it does not correct for them. The Final BCD products for Ge (70 and 160 µm) data similarly include the following items: fully calibrated slope images, the errors associated with the slope images, and mask images.

Whether post-BCD products are sufficient for a particular program depends upon the specific scientific needs of that program. For example, the SSC MIPS pipeline does not provide the type of refined data processing needed to achieve "super-resolution" imaging that can be achieved with the available MIPS data acquisition modes. Observers requiring this type of processing in the combining of BCD images will need to apply standard techniques, or develop new ones as required by the demands of the data and/or the science goals. The software used for the Post-BCD pipeline is available (along with copious documentation) on the SSC website.

The MIPS post-BCD products are derived per AOR, and do not extend to multiple AORs in an Observer's program. It is composed of FITS image files of size $M{\times}N$ (as defined by the specific observation or map) that are the result of optimally mapping the constituent frames into the post-BCD data frame with minimal loss in resolution. There is one post-BCD FITS file per band per selected wavelength per AOR, except for scan map AORs where there are always 3 files per AOR.

The pointing offsets used in mapping of individual frames into the post-BCD products are derived from an optimal offset estimation method, e.g., point source extraction. If insufficient point sources are available, then the data are aligned using spacecraft and CSMM pointing knowledge only. Individual frames are adjusted such that the overlap regions are statistically similar, within the limitations of the large-scale background and slow response characteristics of the Ge:Ga detectors. This implies that some edges might be visible for certain types of varying backgrounds. This will especially be true for adjacent legs of a scan map over a region with significant and varying background. The absolute level to which edge effects and scan leg baseline offsets will be removed in the post-BCD products will not be known with certainty until more flight data are obtained. For observers requiring very uniform maps of extended source regions, there should be some expectation of additional required data processing to correct for edge effects and scan leg baseline offsets in the Ge:Ga data of a scan map. Such effects are minimized in the other MIPS observing modes. The zodiacal

backgrounds are maintained only within an individual AOR.

Observers who wish to produce maps of areas larger than can be imaged in a single AOR could merge either the individual BCDs or the post-BCD mosaics themselves, in order to obtain the combined dataset. Characterizable Ge:Ga long-term detector behavior are included as part of the processing as needed. No background or sky subtraction is made as part of the post-BCD products processing. Ancillary data images similar to those included with Basic Calibrated Data are also provided with the post-BCD products.

Calibration uncertainties from the SSC MIPS data pipelines will decrease with time and are expected to meet the full capabilities of the instrument and the Observatory. For planning purposes, it is of value to observers to understand the initial level of uncertainty expected in the data processing from a calibration standpoint. The pipeline expectations (which include expected initial flat field uncertainties) are as follows:

- Relative point source calibration within same AOR and from AOR to AOR is 5% for 24 µm, 7% for 70 µm, and 12% for 160 µm.
- Emission structures extending over several array sizes (applies primarily to scan mapping), and that fall within a single scan leg, will be preserved to about 10% of relative calibration.
- Absolute calibration will evolve towards lower uncertainty with time. The initial absolute calibration should not be expected to be better than 20% to 30% for Ge.

This level of MIPS pipeline performance is contingent on proper observation design (one that takes full advantage of the standard MIPS calibration), and a suitable target selection. Suitable target selection includes an assessment of the impact of the general background to be observed during the observation, the MIPS saturation limits, and any bright sources that come within the MIPS fields of view.

Pointing uncertainties will improve with time as well. Absolute pointing knowledge of the focal plane array pixel centers based on the Pointing Control System is better than:

- 24 µm : about 1.4″, 1σ radial
- 70 µm : about 1.7″, 1σ radial
- 160 µm : about 3.9″, 1σ radial

From the science data frames themselves, refined pointing knowledge of sources must be based on centroiding of objects visible in a MIPS image. The uncertainties will depend on the signal-to-noise ratio of those source detections and the related uncertainty in the centroids, and on the accuracy of prior knowledge of the positions of some or all of the sources. Uncertainty in relative positions between sources (source offsets) can be much lower than those

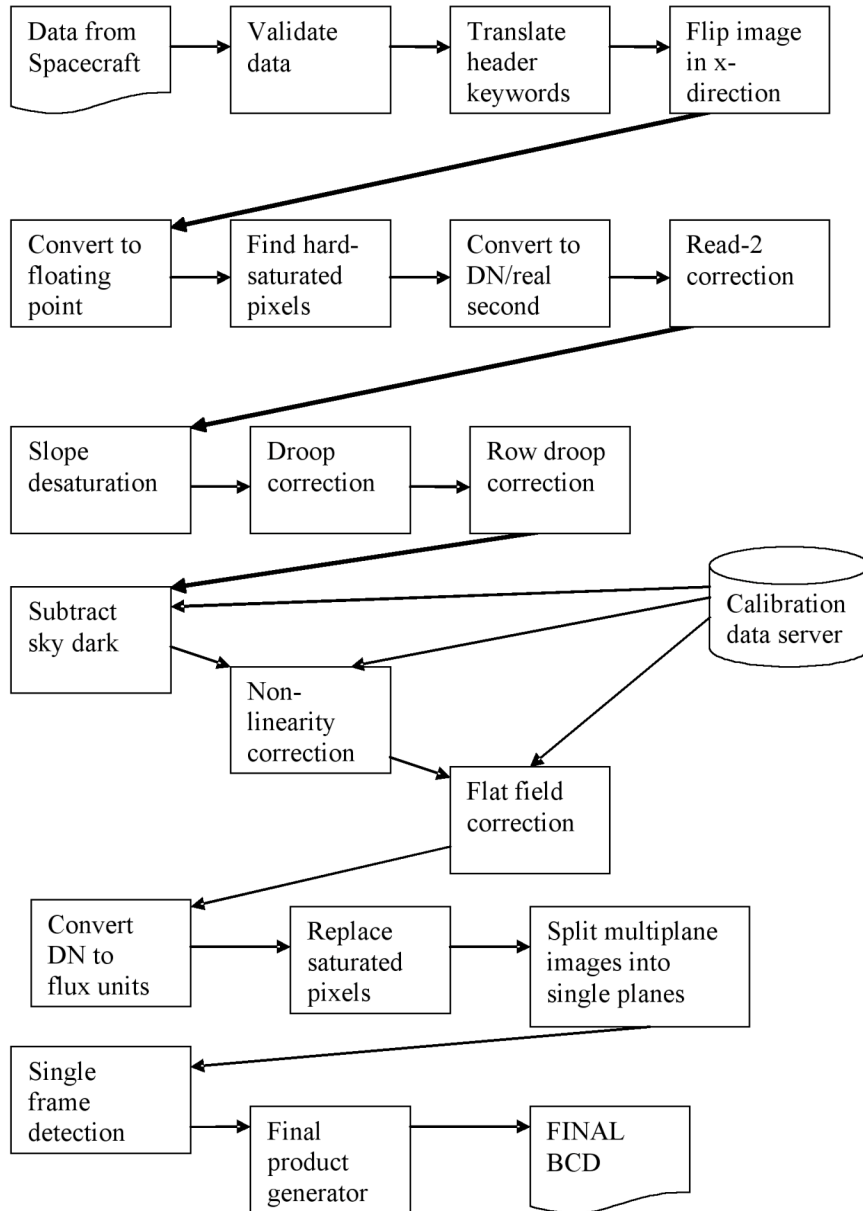associated with absolute pointing knowledge, for detections that can support this type of analysis.

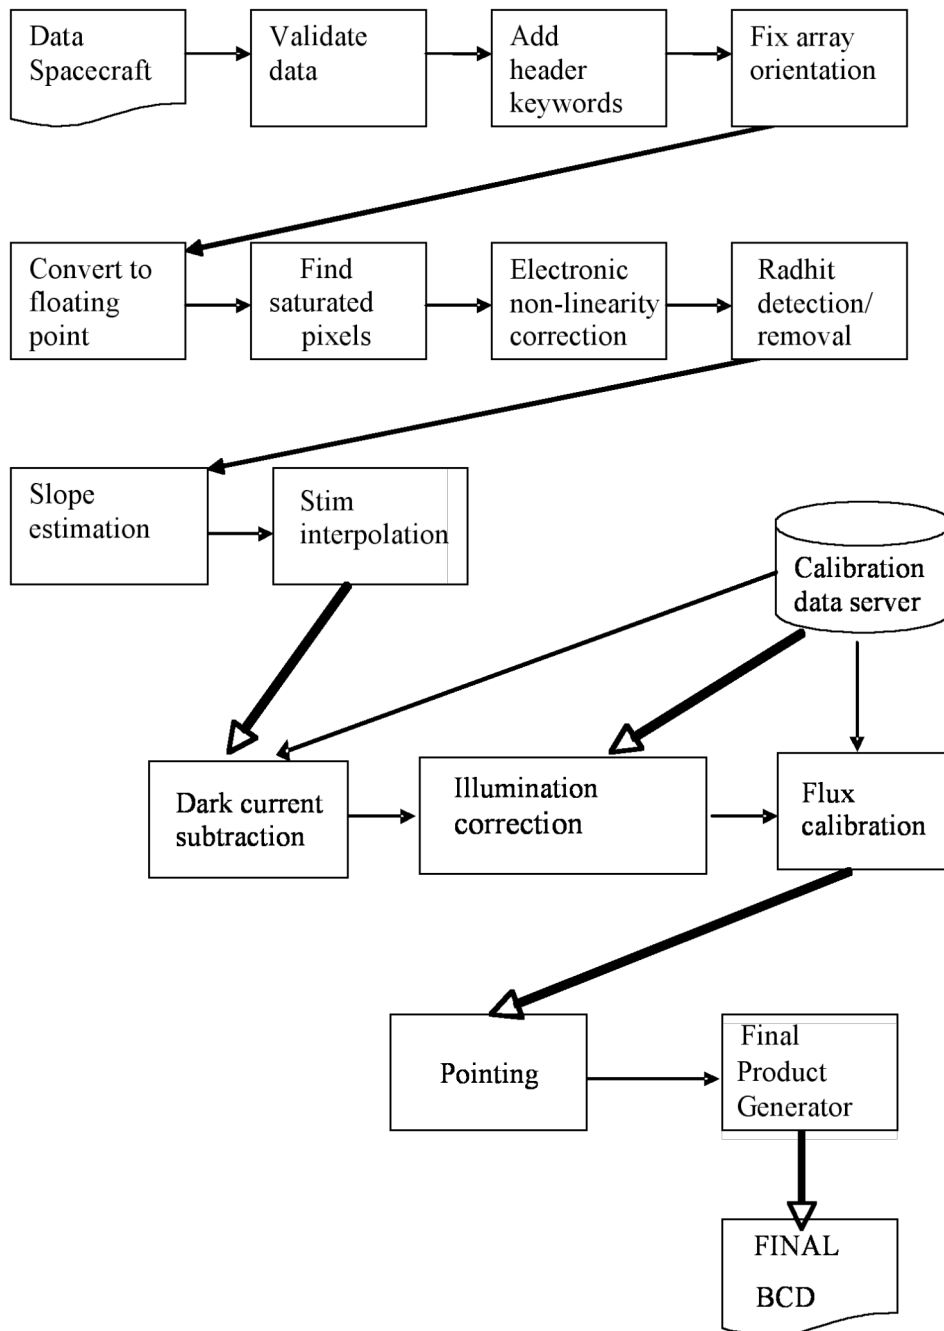**Figure 8.46: MIPS 24 μm science pipeline.**

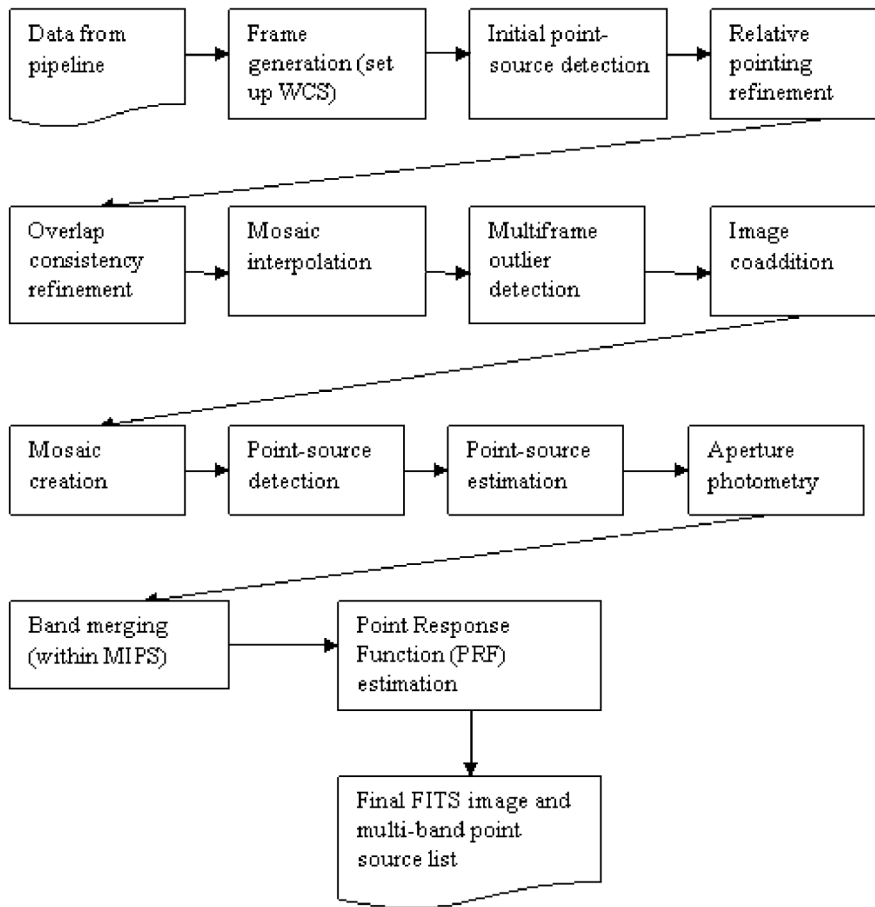**Figure 8.47: MIPS Ge (70 and 160 μm) science pipeline.**

Figure 8.48: The MIPS post-BCD pipeline. Code and extensive documentation are available at the SSC website.

As all types of calibration improve, the archive data will be reprocessed to reflect this evolution. Observers proposing for observations very early in the Spitzer mission need to consider this type of calibration evolution while planning their data processing and analysis needs. However, all Spitzer data will ultimately receive the best available processing and calibration from the pipelines, and the MIPS data archive will have uniform uncertainties.

9 Appendices

9.1 Infrared Flux Units

The infrared flux density from a point source is most commonly given in units of Jansky (Jy) where:

$$1 \text{ Jy} = 10^{-23} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ Hz}^{-1} = 10^{-26} \text{ Watts m}^{-2} \text{ Hz}^{-1} = F_\nu \qquad (10.1)$$

The conversion between Janskys and flux density in W m^{-2} per unit wavelength is given by:

$$F_\nu \times 10^{-26} \times c / \lambda^2 = F_\lambda \qquad (10.2)$$

The infrared flux density from an area on the sky, such as the surface brightness of an extended object, or the background emission, is commonly given in 10^6 Jy steradian^{-1} = 1 MJy sr^{-1}. Another common unit is Jy per square arcsecond; 1 MJy sr^{-1} = 2.350443 × 10^{-5} Jy arcsec^{-2}.

Table 9.1−Table 9.4 list the zero magnitude fluxes for various common optical and infrared filters. Note that the photometric system of filter sets can vary, depending on the manufacturer. The magnitude of a source can be converted to a flux density using:

$$F_\nu = F_0 \times 10^{(-m/2.5)}. \qquad (10.3)$$

There is an on-line Javascript tool available from the SSC website that interactively converts Janskys to magnitudes (or any of a variety of flux units) and vice versa.

Table 9.1: 2MASS system zero points

Passband	Effective wavelength (μm)	Zero point (Jy)
J	1.235	1594
H	1.662	1024
Ks	2.159	666.7

References: Table 2 from Cohen, Wheaton, & Megeath 2003, AJ, 126, 1090; see also 2MASS All-Sky data release web document, http://www.ipac.caltech.edu/2mass/releases/allsky/doc/explsup.html.

Table 9.2: Johnson system zero points

Passband	Effective wavelength (μm)	Zero point (Jy)
U	0.36	1823
B	0.44	4130
V	0.55	3781
R	0.71	2941
I	0.97	2635
J	1.25	1603
H	1.60	1075
K	2.22	667
L	3.54	288
M	4.80	170
N	10.6	36
O	21.0	9.4

References: Allen's Astrophysical Quantities, Fourth edition, 2001, Arthur N. Cox (ed.), Springer-Verlag; Campins, Rieke, & Lebofsky 1985, AJ, 90, 896; Rieke, Lebofsky, & Low 1985, AJ, 90, 900.

Table 9.3: UKIRT system zero points

Passband	Effective wavelength (μm)	Zeropoint (Jy)
V	0.5556	3540
I	0.9	2250
J	1.25	1600
H	1.65	1020
K	2.20	657
L	3.45	2.90
L'	3.80	2.52
M	4.8	163
N	10.1	39.8
Q	20.0	10.4

References: UKIRT web page:
http://www.jach.hawaii.edu/JACpublic/UKIRT/astronomy/conver.html

Table 9.4: MIRLIN zero points

Passband	Effective wavelength (μm)	Zeropoint (Jy)
K	2.2	650.0
M	4.68	165.0
N	10.79	33.4
Qs	17.90	12.4
N0	7.91	60.9
N1	8.81	49.4
N2	9.69	41.1
N3	10.27	36.7
N4	11.70	28.5
N5	12.49	25.1
Q0	17.20	13.4
Q1	17.93	12.3
Q2	18.64	11.4
Q3	20.81	9.2
Q4	22.81	7.7
Q5	24.48	6.7

References: MIRLIN web page (http://cougar.jpl.nasa.gov/MIRLIN/filters.html)

9.2 Infrared Backgrounds

Various astronomical sources emit radiation in the infrared part of the spectrum. Cool stars (M class) have their peak emission just short of the near infrared. However, stars with dusty envelopes or shells and circumstellar disks can be quite bright in the infrared. Regions of star formation, HII regions, and planetary nebulae are strong infrared sources. The (relatively) cool interstellar medium in galaxies has an infrared component. There are also ultra-luminous infrared bright galaxies that are very strong sources of infrared radiation.

As in the optical, the infrared zody is concentrated toward the ecliptic with weaker emission, by approximately a factor of 4, toward the ecliptic poles. The infrared zody is strongest from about 5 μm to about 30 μm with peak emission at about 10 μm. The infrared zody has structure on most scales and, as observed from Earth, varies from season to season. The intensity of the infrared zody will also vary with solar elongation, or how close to the Sun one is pointed. The infrared zody is difficult to model.

As one moves to longer wavelengths (≥ 100 μm), diffuse Galactic emission from dust clouds in the interstellar medium becomes the dominant contribution to the infrared background. This infrared cirrus is patchy, with higher concentrations found in the Galactic disk and toward the Galactic center. However, it is

important to realize that the cirrus is ubiquitous, and it is critical to examine the IRAS maps or radio maps of cirrus tracers when planning longer-wavelength observing. Far-infrared emission from external galaxies in the field of view will add to the overall background flux.

For more information on infrared backgrounds, see the "Backgrounds" section of the Infrared Compendium, available on the SSC web pages.

9.3 Solar System Objects Included in Bright Object Avoidance

Also see the list of bright inertial objects on the SSC website.

- Earth
- Moon
- Mars
- Jupiter
- Saturn
- Uranus
- Neptune
- 4 Vesta
- 6 Hebe
- 1 Ceres
- 7 Iris
- 15 Eunomia

10 Acronyms/Glossary

Note: The following lists the acronyms used in this manual. A larger list of acronyms is available on the SSC website, http://ssc.spitzer.caltech.edu/.

Table 10.1: Acronyms used in this manual

ADC	Analog to Digital Converter
ADU	Analog to Digital Units
AOR	Astronomical Observation Request(s)
AOT	Astronomical Observation Template(s)
BCD	Basic Calibrated Data
BIB	Blocked Impurity Band
C&DH	Command and Data Handling
CA	Cryogenic Assembly
CE	Combined Electronics
CIB	Cold Interface Board
CMOS	Complementary Metal-Oxide Semiconductor
CP	Call for Proposals
CR	Cosmic Ray
CSE	Control Section Electronics
CSFS	Control Section Flight Software
CSMM	Cryogenic Scan Mirror Mechanism (MIPS scan mirror)
CTA	Cryogenic Telescope Assembly
CTIA	Capacitive TransImpedence Amplifier
DAC	Digital to Analog Converter
DCE	Data Collection Event
DCS	Double Correlated Sampling
Dec	Declination
DN	Data Number
DQE	Detector Quantum Efficiency
DSN	Deep Space Network
EDAC	Error Detection and Correction
EEPROM	Electrically Erasable Programmable Read Only Memory
FAQ	Frequently Asked Question
FITS	Flexible Image Transport System
FLS	First Look Survey
FOV	Field-of-View
FPA	Focal Plane Array
FPMA	Focal Plane Mount Array
FWHM	Full Width at Half Maximum
GSFC	Goddard Space Flight Center
GO	General Observer(s)
GTO	Guaranteed Time Observer(s)
HGA	High Gain Antenna
HST	Hubble Space Telescope
IBC	Impurity Band Conduction

ICRS	International Celestial Reference System
IER	Instrument Engineering Request – like an AOR, but for specific engineering tasks not able to be accomplished using an AOR.
IOC	In-Orbit Check out
IPAC	Infrared Processing and Analysis Center
IR	Infrared
IRAC	InfraRed Array Camera
IRAS	InfraRed Astronomical Satellite
IRS	InfraRed Spectrograph
IRU	Inertial Reference Unit
ISO	Infrared Space Observatory
ISSA	IRAS Sky Survey Atlas
Jy	Jansky
LGA	Low Gain Antennae
LIRG	Luminous InfraRed Galaxy
MIC	Multi-Instrument Chamber
MIPS	Multiband Imaging Photometer for Spitzer
mJy	milliJansky
MJy	megaJansky
MPC	Minor Planet Center
MUX	Multiplexer
OPZ	Operational Pointing Zone
PAO	Period of Autonomous Operations
PCRS	Pointing Control Reference Sensor
PCS	Pointing & Control System
PIN	Positive Intrinsic Negative (Diode)
PM	Primary Mirror
PRF	Point Response Function (like a PSF, but having been transmitted through the telescope+instrument optics)
PSF	Point Spread Function
PSSC	Point Source Sensitivity - Continuum
PSSL	Point Source Sensitivity - Line
PST	Point Source Transmission
PUI	Peak-Up Imaging
RA	Right Ascension
RAM	Random Access Memory
RCS	Reaction Control System
rms	Root Mean Square
SAO	Smithsonian Astrophysical Observatory
SBRC	Santa Barbara Research Center
S/C	SpaceCraft
SED	Spectral Energy Distribution
SIRTF	Space InfraRed Telescope Facility, Spitzer's old name
SLO	Second-Look Observation
SM	Secondary Mirror

SODB	Science Operations Database
SOM	Spitzer Observer's Manual
S/N	Signal to Noise
Spot	Previously listed here as an acronym meaning "Spitzer Planning Observations Tool" but now it is simply a proper noun
sr	Steradians
SSC	Spitzer Science Center
SSC website	http://ssc.spitzer.caltech.edu/
SSO	Solar System Objects/Observations
ST	Star Tracker
ST/IRU	Star Tracker/Inertial Reference Unit
TAC	Time Allocation Committee
ToO	Target of Opportunity
TPG	Time Pattern Generator
TP, TPM	Total Power Mode
QE	Quantum Efficiency
WASS	Wide Angle Sun Sensor
WEA	Warm Electronics Assembly

11 Index

2

2MASS, 14, 36, 81, 115, 126, 131, 195, 202, 213, 332

A

anneal, 51, 81, 116, 143, 144, 254, 259, 261, 299, 323

aperture, 8, 9, 30, 38, 39, 40, 43, 45, 65, 96, 109, 145, 148, 149, 150, 157, 182, 189, 204, 235, 248, 294, 297

archive, 6, 7, 9, 13, 14, 55, 123, 126, 127, 228, 302, 324, 330

array coordinates, 107, 110, 118, 120, 121, 122, 182, 219, 281, 313, 314

Astronomical Observation Request (AOR), 12, 24, 34, 51, 57, 58, 61, 65, 66, 89, 105, 108, 115, 126, 129, 139, 168, 169, 211, 228, 268, 271, 303, 306, 307, 314, 321, 324

Astronomical Observation Template (AOT), 9, 12, 50, 51, 98, 105, 115, 168, 171, 187, 211, 238, 265, 283, 307, 321

B

Baseline Instrument Campaign (BIC), 31, 57, 65

Basic Calibrated Data (BCD), 14, 76, 126, 127, 129, 228, 232, 306, 324, 326, 327

best observing practices
IRAC, 115
IRS, 204
MIPS, 309
MIPS photometry, 313
MIPS scan, 319
MIPS SED, 320
Solar System objects, 63

boresight, 24, 33, 34, 35, 43, 226, 268

bright object, 24, 25, 30, 31, 65, 115, 118, 180, 204, 269, 319, 335

C

calibration, 9, 15, 19, 31, 33, 34, 50, 58, 62, 68, 80, 86, 94, 109, 112, 115, 123, 127, 129, 155, 169, 171, 179, 200, 220, 226, 230, 232, 253, 254, 255, 256, 259, 261, 263, 269, 272, 284, 299, 302, 303, 313, 321, 322, 324, 327

celestial coordinates, 66, 107, 120, 122, 169, 182

cirrus, 54, 189, 194, 202, 294, 296, 334

color correction, 220, 226

Common Electronics (CE), 138, 140, 259

confusion, 7, 10, 41, 100, 101, 120, 189, 258, 294, 309

constraints, 12, 18, 24, 30, 33, 35, 52, 56, 58, 62, 120, 121, 194, 212, 217, 255, 262, 268, 303, 314, 315
chain, 59
follow-on, 60
group-within, 59
sequence, 59
shadow, 60, 62, 295
timing, 30, 59, 62, 64

coordinate system, 19, 25, 33, 43, 71

cosmic rays, 86, 121, 127, 130, 142, 143, 155, 171, 173, 192, 232, 240, 263, 290, 295, 324

cross-scan, 267, 268, 270, 273, 279, 282, 287, 290, 311, 313, 316

cryogens, 6, 8, 18, 24, 38, 39, 136, 239

cryostat, 18, 21, 41, 43

D

dark, 41, 84, 85, 96, 100, 125, 127, 129, 130, 143, 155, 156, 157, 173, 179, 193, 226, 231, 232, 241, 246, 250, 258, 293, 322, 323, 325

Data Collection Event (DCE), 127, 129, 140, 181, 199, 228, 299, 324

Data Handbook, 2, 3, 15, 71, 77, 92, 124, 126, 129, 179, 228, 253, 258, 269, 298, 302, 307, 320, 322, 325

Director's Discretionary Time (DDT), 6, 7, 13, 65, 66

distortion, 32, 71, 73, 115, 125, 128, 131, 177, 180, 226, 231, 233, 271, 275, 306, 321, 324

dither, 34, 35, 53, 76, 82, 83, 86, 97, 105, 106, 109, 110, 113, 115, 122, 179, 180, 183, 219, 222, 226, 238, 273, 282, 290, 303, 305, 308, 310, 311, 312, 313, 324

Double Correlated Sampling (DCS), 173, 178, 192, 228

droop, 156, 204, 225, 232, 246, 253, 293, 325

E

ephemeris, 14, 35, 61, 64, 170, 198, 204

extended source, 51, 54, 61, 96, 102, 125, 156, 157, 165, 174, 182, 183, 191, 192, 201, 202, 205, 215, 227, 238, 258, 265, 271, 290, 298, 299, 300, 301, 310, 316, 326, 332

F

fault protection, 33, 34, 94, 95

filter, 69, 71, 77, 130, 145, 147, 177, 178, 181, 187, 191, 193, 194, 195, 201, 213, 226, 239, 242, 245, 258, 298, 307, 332

filter transmission, 78, 187. *See also* spectral response

first frame effect, 89, 106, 117

First Look Survey (FLS), 7

FITS format, 15, 126, 132, 228, 321, 324

flat, 86, 97, 99, 110, 112, 124, 127, 130, 155, 173, 179, 193, 221, 227, 228, 231, 254, 322, 323, 325

focal plane, 11, 30, 44

Fowler sampling, 89, 92, 93, 98, 106, 129