

Processing for the MIPS GAL 24 μm Survey of the Inner Galactic Plane

D. R. MIZUNO,¹ S. J. CAREY,² A. NORIEGA-CRESPO,² R. PALADINI,² D. PADGETT,² S. SHENOY,²
T. A. KUCHAR,¹ K. E. KRAEMER,³ AND S. D. PRICE³

Received 2008 February 28; accepted 2008 July 10; published 2008 August 25

ABSTRACT. We describe improvements to the data processing pipeline for the *Spitzer* MIPS GAL survey of the Galactic plane. These involve both improvements to the processing of the 24 μm data, in particular the handling of saturations and near-saturated data, and mitigation of various artifacts not corrected in the basic calibration pipeline. The artifacts addressed postpipeline are typically caused by passage across bright point sources very common in the Galactic plane, and include column-to-column “jailbar” striping, latency effects resulting in both short-duration afterimages and long-duration responsivity depressions, scattered light, and background-level mismatches. We describe the artifacts phenomenologically and present in detail the algorithms developed to correct them.

1. INTRODUCTION

MIPSGAL is a survey of the inner Galactic plane at 24 and 70 μm using the Multiband Imaging Photometer System (MIPS) aboard the *Spitzer Space Telescope*. Rieke et al. (2004) provide a description of the MIPS instrument, and Werner et al. (2004) describe the *Spitzer* spacecraft. The scientific objectives, observing strategy, data products, and early results of the MIPSGAL project are described elsewhere (Carey, S. J., et al. 2008, in preparation).

MIPSGAL comprises two separate observing projects: MIPSGAL I (program ID 20597) covered Galactic longitudes $l = 10$ to 65° , and $l = 295$ to 350° , to Galactic latitude $|b| \leq 1^\circ$. MIPSGAL II (program ID 30594) extended the survey to higher latitudes, $|b| \leq 3^\circ$ in the Galactic center region, $l = 350$ to $l = 10^\circ$, excluding a $9^\circ \times 2^\circ$ region aligned on the Galactic center, which was covered in a different project (program ID 20414, principle investigator F. Yusef-Zadah). For survey completeness, however, this latter data set has been processed with the reduction procedures developed for MIPSGAL and is incorporated into the MIPSGAL data products.

The image products are a set of 353 $1.1^\circ \times 1.1^\circ$ mosaics, with a pixel size of $1.25''$ and a resolution of $5.6''$. The mosaics are centered at integral Galactic longitudes and half-integral latitudes, and are made in tangent plane projections. In addition to the sky mosaics, corresponding images of coverage, uncertainty, and data flags are provided. See Carey, S. J., et al. (2008, in preparation) for a complete description. The mosaics are

available from the Infrared Processing and Analysis Center (IPAC).⁴

This paper details the steps involved in processing the raw 24 μm data into the final mosaics. Section 2 provides a brief outline. The 70 μm data set requires very different processing, due primarily to the responsivity variations in the Ge:Ga arrays, and is the subject of a separate paper (Paladini R., et al. 2008, in preparation).

The processing presented here falls into two general categories: calibration and artifact mitigation. A comprehensive discussion of the standard calibration steps is outside the scope of this paper (see Gordon et al. 2005), but we have made modifications to the Basic Calibrated Data (BCD) pipeline, in particular improvements to the linearity correction, droop correction, and handling of saturated pixels. These are presented in § 3.

The 24 μm data set suffers from a number of residual artifacts that are not addressed in the basic calibration, among them latency effects, in which passage across a bright source causes both short-term bright afterimages and a long-term responsivity depression, column-by-column “jailbar” striping, and scattered light on the array due to point sources just out of the field of view. Most of these artifacts are caused by bright sources passing across or near the Si:As array, and the high density of bright sources in the inner Galactic plane makes these artifacts ubiquitous.

The MIPS Data Handbook⁵ describes mitigation techniques for some of these artifacts, but these techniques are either inadequate for the highly structured backgrounds in the MIPSGAL data or are insufficiently effective, and have not been employed in the current processing.

¹Boston College, Institute for Scientific Research, Chestnut Hill, MA, 02467-3800; afrl.rvb.pa@hanscom.af.mil

²*Spitzer* Science Center, California Institute of Technology, Pasadena, CA, 91125; carey@ipac.caltech.edu, alberto@ipac.caltech.edu, paladini@ipac.caltech.edu, dlp@ipac.caltech.edu

³Air Force Research Laboratory, Hanscom AFB, MA, 01731-3010; afrl.rvb.pa@hanscom.af.mil

⁴ Mosaics are available at <http://irsa.ipac.caltech.edu/data/SPITZER/MIPSGAL>.

⁵The MIPS Data Handbook is available at <http://ssc.spitzer.caltech.edu/mips/dh/>.

Instead, we have developed artifact mitigation methods tailored for the MIPS GAL data set. The large amount of data in the MIPS GAL survey necessitates that the correction procedures for these artifacts be as automated as possible. These procedures are detailed in § 4.

Section 5 explains a procedure for correcting background-level discrepancies among the approximately half-million data frames in the survey by minimizing the differences in the frame-to-frame overlaps, so that the resulting mosaics have consistent levels. Section 6 addresses residual errors in the artifact corrections, and § 7 briefly discusses the effectiveness of the corrections and provides a summary.

While the artifact corrections have been tailored to the conditions in the Galactic plane, the correction procedures, in particular those for the jailbars and latency effects, are very general and can be applied to other MIPS data sets. The background-level correction mechanism, in addition, can be applied to other space-based data; the *Spitzer* Infrared Array Camera (IRAC), for example, needs a similar correction.

2. OBSERVATION AND PROCESSING OVERVIEW

The MIPS 24 μm instrument is based on a 128×128 array of Si:As detectors. The field of view of each pixel is $2.55''$ on a side, giving a $5.4'$ field for the array. For the MIPS GAL project, 3 s integrations were used for each frame of data. Transmitted to the ground for each frame are a slope and first-difference image (explained in § 3.1 below). The Basic Calibrated Data (BCD) pipeline converts these images into a single frame of calibrated data (hereafter BCD frame or just BCD), along with a corresponding uncertainty image, and a BCD mask image that uses bit planes to store per-pixel flags such as saturation, likely latency effects, and other conditions.

Observations are performed in units called Astronomical Observation Requests (AORs), which specify integration, pointing, and scanning parameters for the telescope. For MIPS GAL, the fast scan-mapping mode was used (see the *Spitzer* Observer's Manual⁶ for details), tiling the survey region with scans 4 to 6 degrees in length oriented obliquely to the Galactic plane at an angle of $\sim 66^\circ$ east of Galactic north-south. Each AOR consisted of 3 or 4 scan legs. Each line-of-sight in the survey is covered at least 5 times in each of 2 scan legs so that each sky position is imaged in at least 10 BCDs. See Carey, S. J., et al. (2008, in preparation) for a description of the observing strategy. MIPS GAL I and II comprise 493 AORs, each with ~ 1000 BCD frames.

The artifact mitigation is also handled by AORs, in part for convenience and also for the reason that some artifacts are characteristic of the data over complete AORs. The BCDs for each AOR are passed through a sequence of correction procedures,

each mitigating a particular artifact (§ 4). Following the artifact corrections, the overall levels of the BCDs are matched in overlapping regions to ensure background consistency (§ 5). The corrected, overlap-matched BCDs are then combined into the mosaics using the MOPEX⁷ software provided by the *Spitzer* Science Center. The products are the basic sky brightness image, coverage and uncertainty images, and a mask image with an accounting of selected mask flags that affect each mosaic pixel.

The mosaics are inspected visually, with residual artifacts and other problems recorded, and corresponding remedial actions are taken. The corrections are then reapplied, and the mosaics regenerated and inspected. The process is repeated as necessary.

3. MIPS GAL BCD PIPELINE

The *Spitzer* Science Center (SSC) version S14.4 of the BCD pipeline was available at the time the MIPS GAL data were processed, but the pipeline corrections for linearity and substitution of the first-difference pixel values (described in § 3.2) were not optimal, which resulted in the intensities above 300 MJy sr^{-1} in that version being unreliable. Instead, in consultation with the MIPS Instrument Support Team at the SSC, we developed an offline version of the BCD pipeline that provides these corrections using the raw data and calibration files available from Leopard, the standard software tool for retrieving data from the *Spitzer* archive.⁸

Our modified pipeline starts with the raw data, uses the SSC calibration files, and performs all the basic calibration steps of their pipeline: identification of saturated pixels, “read-2” correction, row-droop and droop correction, dark subtraction, electronic linearization, flat-fielding, conversion to physical units, and pixel replacement. Gordon et al. (2005) provide a description of these steps. The linearity solution implemented in our pipeline is identical to that in SSC pipeline versions S16.1 and later, and in most other respects, our pipeline is identical to the SSC pipeline. We have, however, made several modifications appropriate to bright, saturation-filled regions of the sky. We have tailored the identification of saturated pixels and the droop correction to better account for pixels that are saturated in the first difference.

3.1. Data Collection Overview

In the fast scan-mapping mode used for MIPS GAL, the integrations are for 2.6 s. The accumulating signal in a pixel (data number or DN) is sampled nondestructively 6 times at $\Delta t = 0.52$ s intervals, a mode labeled “sampling up the ramp” (SUR) in MIPS nomenclature. Let DN_0 through DN_5 represent the six sampled values, or “reads.” To save bandwidth, these raw values are not transmitted to the ground but are processed

⁶The *Spitzer* Observer's Manual is available at <http://ssc.spitzer.caltech.edu/documents/som/>.

⁷The MOPEX software is available for download at <http://ssc.spitzer.caltech.edu/postbcd/mopex.html>.

⁸The Leopard user's guide is available at <http://ssc.spitzer.caltech.edu/documents/leopard/>.

onboard the spacecraft into two derived quantities (for this processing, the initial sample, DN_0 , is discarded due to electronic artifacts, so the effective integration period is about 2 s). The *slope* is calculated by linear regression over the five included samples DN_1 to DN_5 , and the *first-difference* is the difference of the first two. These values are measures of the signal accumulation rate and hence the incident brightness on the pixel.

The onboard slope calculation accounts for differential errors and correlations between the DN (Gordon et al. 2005), but for illustrative purposes here we show a simplified form. For fitting a linear function to five evenly spaced and equally weighted data points, we have

$$\text{slope} \approx \frac{-2 \times DN_1 - DN_2 + DN_4 + 2 \times DN_5}{10} \quad (1)$$

and, to be explicit,

$$\text{first-difference} = DN_2 - DN_1. \quad (2)$$

Note that the (simplified) slope calculation is independent of DN_3 . The units of both the first-difference and the slope values are DN per sampling interval, or DN per read.

The maximum value of DN is the value at which the signal saturates the analog-to-digital converter, $DN_{\max} = 65,535$. If all samples for a pixel are below DN_{\max} , then both the slope and the first-difference will be accurate measures of the pixel brightness, although the first-difference will of course be noisier. At higher signal levels, DN_5 first begins to saturate and is then followed in turn by the earlier samples. If any of DN_5 down to DN_3 are saturated (i.e., values fixed at DN_{\max}), then the slope will be an underestimate but the first-difference will still be valid; this extends the dynamic range. If DN_2 is saturated, then the first-difference will also be an underestimate; we label this condition a *soft saturation*. If DN_1 is saturated, then both the slope and the first-difference return zero; this is a *hard saturation*.

For the MIPS GAL BCD pipeline, at issue is the criteria for replacing the slope image pixels with the first-difference pixels, and identifying conditions for which the first-difference is in soft saturation. Figure 1 displays the slope value versus the first-difference value for a sampling of data pixels in a selected MIPS GAL AOR centered at about $l = 15^\circ$. The slope and the first-difference are equivalent (to within nonlinearity effects) up to about point A in the figure. At this point, DN_5 first becomes saturated, so the slope beyond this is an underestimate. From equation (1), constant DN_5 (at DN_{\max}) results in a nearly constant slope ($DN_i \propto i$, approximately), and saturation of DN_4 at point B causes a decreasing slope with increasing brightness.

Saturation of DN_3 has no effect on the trending as it does not enter the slope calculation. DN_2 becomes saturated at point C in the figure. After this point—the onset of soft saturation—with increasing brightness DN_1 continues to increase but $DN_2 = DN_{\max}$, so the first-difference value decreases to zero, when DN_1 also saturates. In this regime, the slope and the first-

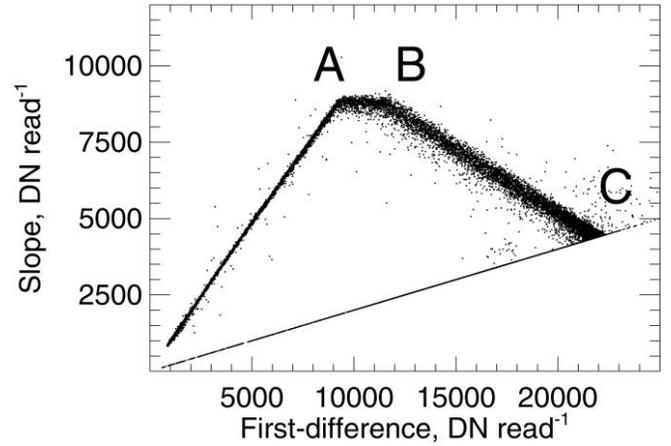


FIG. 1.—Sampling of slope vs. first-difference values from a MIPS GAL AOR. Increasing incident brightness proceeds clockwise from the origin around the triangular shape. A, B, and C indicate the onset of saturation for read samples DN_5 , DN_4 , and DN_2 , respectively.

difference are functions only of DN_1 (and DN_{\max}), so there is a fixed relation between the two values. Equations (1) and (2) show that the first-difference/slope ratio is identically 5 in this situation, so with increasing true brightness the pixel loci descend along a straight line to hard saturation at the origin.

3.2. Pixel Replacement

The first criterion for replacing the slope image pixels with the first-difference pixels is the point at which the final read, DN_5 , first saturates; as we have seen this is at point A in Figure 1, or at a first-difference value of about 9000 DN per read. We see also from the plot that all first-differences above this value correspond to a slope that is beyond the threshold DN_5 saturation point, although the first-difference values themselves might be in soft saturation. Thus, using a threshold first-difference value is a reliable criterion for identifying saturated slope values.

The conversion to intensity is $1 \text{ DN s}^{-1} = 4.54 \times 10^{-2} \text{ MJy sr}^{-1}$ (Engelbracht et al. 2007) or $1 \text{ DN per read} = 8.73 \times 10^{-2} \text{ MJy sr}^{-1}$. A threshold of 500 MJy sr^{-1} in the first-difference image is selected as the pixel replacement criterion. This corresponds to $\sim 6000 \text{ DN per read}$ and so is conservatively well short of saturation in the slope image.

The second criterion is to identify and flag pixels in the soft-saturation regime, i.e., the pixels lying along the straight line from point C in the figure to the origin, for which DN_2 is saturated, and the first-difference is no longer a valid brightness measure. The first-difference value at point C, about 22,000 DN per read, is still valid, and this corresponds to about 1900 MJy sr^{-1} . This is roughly the upper limit pixel value in the MIPS GAL images, although values up to $\sim 2300 \text{ MJy sr}^{-1}$ are seen.

For pixels in soft-saturation, the ratio of the first-difference to the slope measurement is exactly 5. For pixels below the

soft-saturation threshold brightness, the ratio is always smaller than this, monotonically increasing from ~ 1 for unsaturated data up to the soft-saturated value. We use this result to identify soft-saturated pixels, defining soft saturation as any pixel for which the ratio is within 10% of the exact value, i.e., if

$$\frac{\text{first-difference}}{\text{slope}} \geq 4.5, \quad (3)$$

the pixel is considered soft saturated, and is both flagged in the BCD mask and replaced with the Institute of Electrical and Electronics Engineers (IEEE) Not-a-Number (NaN) value in the BCD image. Hard-saturated pixels are also flagged and replaced with the NaN value.

3.3. Droop Correction

The droop is an effect in which the entire array is elevated by an offset, with a magnitude depending on the total signal on the array. The correction is to total up the pixel values over the array, scale appropriately, and subtract. In the SSC pipeline, the droop offset is calculated as 0.248 times the mean pixel brightness over the array (Gordon et al. 2005), and that coefficient is carried through in our pipeline.

The presence of saturated pixels, however, lowers the measured total from the true total signal as these pixels always have lower values than the true brightness. With the standard pipeline, we see the result as elevated background levels in BCDs containing bright saturated sources.

For soft saturations, the true brightness can be estimated crudely with the knowledge of what happens at the onset of saturation for DN_1 and DN_2 : When DN_2 just saturates, it means that DN_{\max} is reached after three sample periods (integration starts at about one sample period prior to DN_0). Thus, the true accumulation rate is $DN_{\max}/3 \approx 21,800$ per read, and we see this in Figure 1; point C is at a first-difference value of $\sim 22,000$ DN per read, and is still a valid measure. When DN_1 just saturates, the true data rate is likewise about $DN_{\max}/2 = 32,768$ per read, while the first-difference value drops to zero. The measured values of the first-difference in the soft-saturation regime thus fall in the interval $(DN_{\max}/3, 0)$ and these values map into an estimated true accumulation rate of $(DN_{\max}/3, DN_{\max}/2)$. Assuming a linear mapping, we have

$$\text{true-rate} \approx \frac{DN_{\max} - \text{first-difference}}{2} \quad (4)$$

in units of DN per read.

For the purpose of calculating the total signal in the droop correction, the soft-saturated pixels are replaced with this estimate of the true brightness. Hard-saturated pixels are replaced with $DN_{\max}/2$ per read. These substitutions yield a significant improvement to the correction; a residual droop is typically seen only in BCDs containing either extremely bright point sources

with large hard-saturated cores, or large regions of hard-saturated extended emission.

4. ARTIFACT MITIGATION

Several significant artifacts remain after the pipeline production of the basic calibrated data, most of which are induced by bright sources in or near the BCD frames. The large number of bright point sources in the MIPSGAL data set precludes a solution such as simply deleting affected frames. The objective is to correct the artifacts to whatever extent possible, and this effort has been successful in removing the jailbar pattern and bright and dark latency effects from the data. Other artifacts, such as those due to stray light and internal reflections, are too irregular to correct, and the solution is to flag the affected regions of the BCDs in the masks and subsequently omit them from the final mosaics.

In this section, we describe each artifact and present the mitigation procedure we have developed. A brief discussion of the effectiveness and limitations of the procedures is presented in § 7.

4.1. Jailbars

A very pronounced, regular striping pattern commonly appears in images that contain bright sources. An example is shown in Figure 2. The striping consists of column-to-column variations of levels, and the pattern repeats every four columns, thus inspiring the label ‘‘jailbars.’’ The deviations are typically a few MJy sr^{-1} but can be tens of MJy sr^{-1} in the worst examples. The pattern is also typically different above and below the

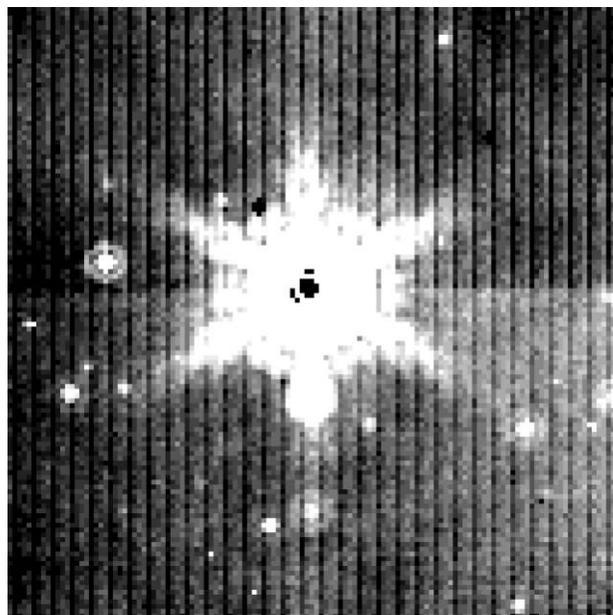


FIG. 2.—Example of a BCD containing the jailbar artifact. Note that the pattern is similar above and below the triggering source, but differs in intensity.

presumed triggering point source, as shown in the figure. In nearly all cases, the jailbar pattern has a single bright-source trigger and consists of differing above-and-below sections, but a few cases have been observed of two triggering sources and three distinct jailbar sections. The jailbar effect is always a depression of levels in the columns.

No minimum point source flux for causing this effect has been determined, but it is very seldom seen with sources that are not saturated at the center ($\lesssim 4$ Jy) and very common with sources >10 Jy. Large regions of saturated extended emission can also cause jailbars. There is no direct correlation of the brightness of a source and the strength of the jailbar pattern; a given source will usually be scanned over several successive BCDs, and the pattern can be different in all of them. The pattern is often worst when the source is near the top or bottom of the array, but there is otherwise no general rule.

No cause for this effect has been identified, but the pattern is clearly related to how the pixels in the Si:As array are read out. The pixels are read out in four channels, each readout channel consisting of 32 columns of the 128 total. The columns are interleaved such that each readout samples every fourth column in the array. When a very bright source falls on the array, we postulate that the pixels in a given readout are depressed by an amount depending in some way on the fluence on the brightest pixel in the readout. As there will generally be differing brightest pixels in the separate readouts, there will likewise be differing depressions, thus producing a striping pattern in the BCD that is repeated every four columns.

Engelbracht et al. (2007) suggest that the differing above-and-below pattern is a consequence of two jailbar triggers, one a bright but unsaturated point source, or cosmic ray, which affects the whole array, and the other—the obvious saturated source in the image—which alters the pattern following the read of the saturated pixels, i.e., above the source in the array. This would imply, however, that uniform jailbar patterns in the absence of a saturated source should be very common, but in fact are comparatively rare in the MIPS GAL data.

Whatever the origin of the jailbar effect, the mitigation makes use of its phenomenological characteristics. The striping is due to differential depressions of levels over entire readouts, and so the correction is a global adjustment to each of the readouts in a section. The interleaving of the readout columns means that, in a given BCD, the columns of each readout collectively observe approximately the same region of the sky. After removal of any global gradient, the differences in the readout medians are assumed to solely reflect any jailbar errors; residual differences due to true sky brightness variations between the readouts are presumed small enough to ignore.

The steps in the jailbar correction for each BCD are:

1. *Remove flat-field correction.* The jailbar deviations are uniform in a readout *before* the BCD pipeline flat-field correction.
2. *Select section boundaries.* Array rows containing hard or soft saturations are considered potential boundaries (i.e., cross a

triggering source). Contiguous rows containing a saturation are collapsed to a single assumed boundary. The regions between boundaries are the jailbar sections. If there are no saturations, the row containing the peak pixel in the array is considered the boundary. Multiple boundaries are allowed, but a minimum of 5 rows per jailbar section is imposed.

3. *Determine corrections per section.* The readout medians are calculated after removal of any overall gradient in the section. As the jailbar effect is a depression, the peak readout median is considered closest to truth, and the corrections for the other three readouts are the additive offsets necessary to match the peak readout. The depressions are typically up to a few MJy sr^{-1} in magnitude.

4. *Resolve boundaries.* The precise transition point of the jailbar patterns between sections is not known in advance, as saturations usually cover several rows and the transition can occur in any of them. All possible transition points are tested, and the one yielding the lowest variance in the corrected image over the potential transition rows is selected.

5. *Apply corrections, and reapply flat field.*

Note that there is no selection process to distinguish which BCDs suffer from jailbars. All BCDs are passed through the jailbar correction procedure.

To investigate the effect of the jailbar correction on jailbar-free BCDs, the correction procedure was applied to two AORs, one containing quiescent backgrounds (at $l \sim 55^\circ$) and one with bright, structured backgrounds (at $l \sim 8^\circ$). The data arrays were rotated by 90° to ensure that no jailbars would be present in the BCDs (the presence of horizontal stripes should have no effect on the jailbar correction). The resulting jailbar “corrections” are then a measure of the error in the correction mechanism.

The threshold for visual detection of jailbar effects in a given BCD is a readout-to-readout rms of $\sim 0.1 \text{ MJy sr}^{-1}$. For the quiescent AOR, the median rms in the resulting readout corrections was about 0.02 MJy sr^{-1} , and the worst rms was half the detection threshold. For the brighter AOR, the median correction rms was the same, but about 1% of the BCDs fell between one and 3 times the threshold. In all these cases, however, the BCD contained extremely bright structured emission so that the added jailbar pattern was not visually detectable.

Also, the jailbar correction added a very small overall bias to the BCDs, from $\sim 0.1 \text{ MJy sr}^{-1}$ for the quiescent BCDs to $\sim 0.2 \text{ MJy sr}^{-1}$ for the bright, structured ones. This is perhaps not surprising as the jailbar correction is always positive. The absolute background levels in the BCDs are not known to this precision—the uncertainty in the Zodiacal model (Kelsall et al. 1998) is $\sim 10\%$ and so this alone gives an uncertainty $\geq 1 \text{ MJy sr}^{-1}$ for MIPS GAL, and we therefore conclude that the jailbar correction is not adding visually detectable artifacts to the data.

For jailbar examples of light to moderate severity, the peak readout is indeed close to the truth (i.e., in such cases, the peak readout is unaffected by the jailbar depression), and the correction generally results in proper levels in the sections both

above and below the triggering source. In more severe examples, however, all the readouts suffer from some level depression, and matching to the peak readout still leaves an overall level error. Moreover, the effect will generally differ above and below the source, and so there will often be a residual discontinuity of levels across the source after the correction, even though the jailbar pattern itself has been removed. Figure 3 shows before-and-after examples of both these cases.

The situation of residual level errors following the jailbar correction are handled in the overlap correction procedure described in § 5.

4.2. Bright Latents

Bright latents are short-duration afterimages that are observed as “ghost” sources following passage of the Si:As array across a bright source. When pixels in the array are exposed to high brightness levels in a given BCD, those array pixels are biased in the immediately subsequent BCD by about 1% of the original observed brightness, $\sim 0.3\%$ in the second BCD that follows, and decreasing quasi-exponentially for each BCD thereafter. While the effect is most visible for bright localized sources, it is assumed that this phenomenon occurs at all incident brightness levels, though smooth extended emission will cause no obvious artifact.

In the MIPS GAL survey, the data were taken in scan legs in which the pointing is shifted $54''$ or about $1/6$ of the array extent

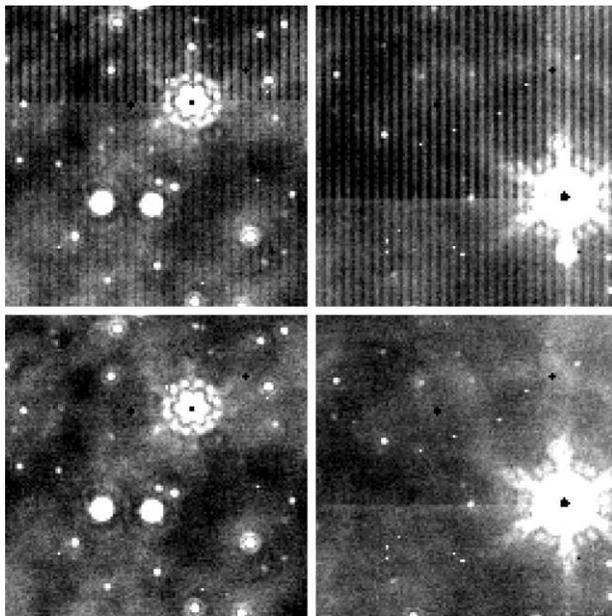


FIG. 3.—Two examples of jailbar artifacts and the results of the correction described in the text. *Left:* The sections above and below the triggering source correct to a consistent background level. *Right:* A discontinuity in level remains after the correction. This residual error is resolved with the overlap correction described in § 5.

between successive BCDs. In crossing a bright source, the source thus strikes the array at five or six successive locations, and the bright latent effect creates a sequence of false pointlike sources, trailing the true source with diminishing brightness. Figure 4 shows an example of bright latents in a BCD.

The effective source flux threshold for producing a barely visible latent is about 100 mJy, for which we see a single faint ghost image following the source. A several-hundred Jansky source will typically produce visible latents for 10–12 subsequent BCDs.

The strategy we take is to model the bright latent response as a function of incident pixel brightness and subtract the modeled latent brightness from downstream BCDs. The function can be sampled directly at low incident brightness values, up to ~ 1500 MJy sr^{-1} , by measuring incident point source peaks and the resulting latent peaks in the BCDs in reasonably quiescent regions. Higher incident levels saturate the array pixels, but the intrinsic peak brightness levels can be determined by measuring the brightness of the fainter outer parts of the source image (Airy disk and/or ring) and scaling the known point response function (PRF) to approximate the total profile. In this way, the response for the first two latent intervals can be characterized empirically.

For incident intensities below a few hundred MJy sr^{-1} , we find a linear relationship between the incident brightness f and the subsequent latent response L , which rolls off at higher incident levels to a maximum latent brightness when the incident brightness is about 18,000 MJy sr^{-1} . The response is apparently constant at higher incident brightnesses. We model

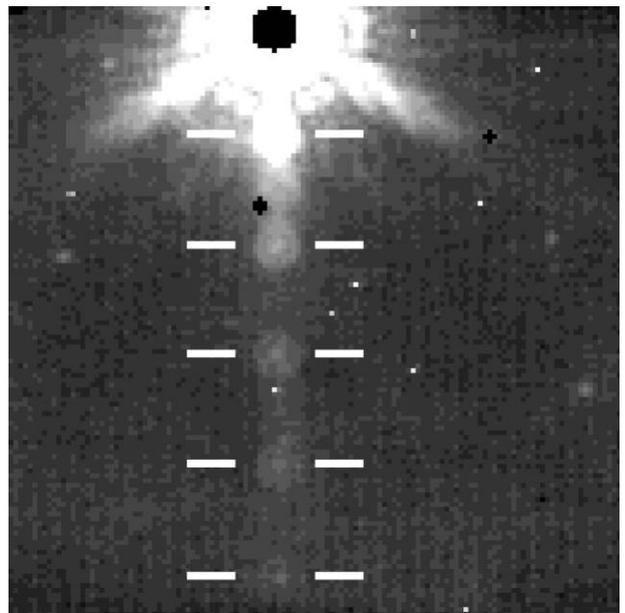


FIG. 4.—Example of a BCD containing bright latents due to passage across the array of a very bright source, ~ 100 Jy. Five latent intervals are visible in this image.

the trending with an exponential plus a linear term:

$$L = A(1 - e^{-\frac{f'}{B}}) + Cf', \tag{5}$$

where

$$f' = \begin{cases} f, & \text{if } f < 18,000 \\ 18,000, & \text{if } f \geq 18,000 \end{cases}, \tag{6}$$

L and f are in MJy sr^{-1} , and the parameters A , B , and C depend on the specific latent interval. For the first two intervals, L_1 and L_2 , the parameters are determined initially by fitting to the empirical data and are then adjusted manually to give improved results in test cases. Table 1 shows the values of A , B , and C for the first two latent intervals.

For the remaining intervals, up to the twelfth latent away from the source, we model the latent brightness by scaling the value of the second latent L_2 . For the first three additional intervals, the scaling was measured from prominent bright latent examples as in Figure 4; the remaining scaling factors were extrapolated from the general trending. Table 1 also shows the scaling factors for these higher-order latents. Figure 5 shows the model latent functions for the first five latent intervals, along with measured incident/latent data points for the first latent interval.

We predict and subtract the latent brightnesses for the successive latent intervals with this model. The basic procedure is to step through the BCDs in a given AOR, and for each given BCD, use the brightness over the whole array to predict the latent effects in subsequent BCDs. In the following BCD, the predicted latent levels are subtracted to give the latent-corrected image, and the process is repeated.

To do this thoroughly, however, we need accurate brightness profiles for all the point sources. Above about 3 Jy, soft saturation renders the center pixel values unusable, and hard saturation sets in somewhat above 4 Jy, so using the data array as-is

TABLE 1
BRIGHT LATENT MODEL PARAMETERS

LATENT INTERVAL	PARAMETERS IN EQUATION (5)		
	A	B	C
1	20.5	2700	0.00048
2	3.4	1600	0.00027
	L_2 scaled by		
3		0.500	
4		0.313	
5		0.200	
6		0.143	
7		0.125	
8		0.111	
9		0.100	
10		0.091	
11		0.083	
12		0.077	

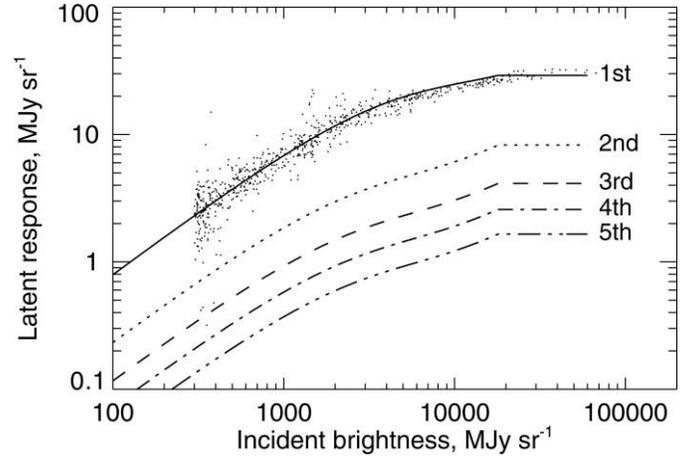


FIG. 5.—Model of the bright latent response as a function of the incident brightness on the array shown for the first five latent intervals. Also shown are measured data points for the first interval.

would significantly undercorrect the latents for these sources, particularly ones much brighter than the saturation threshold. To handle these bright point sources, we create a catalog of positions and fluxes of saturated and near-saturated sources for each AOR. The fluxes of such sources are estimated by performing PRF fitting on the Airy ring of secondary peaks (e.g., Fig. 3 or 6), which themselves become saturated only at fluxes >1500 Jy. In the correction procedure, the true source profile for these saturated sources is modeled by scaling the

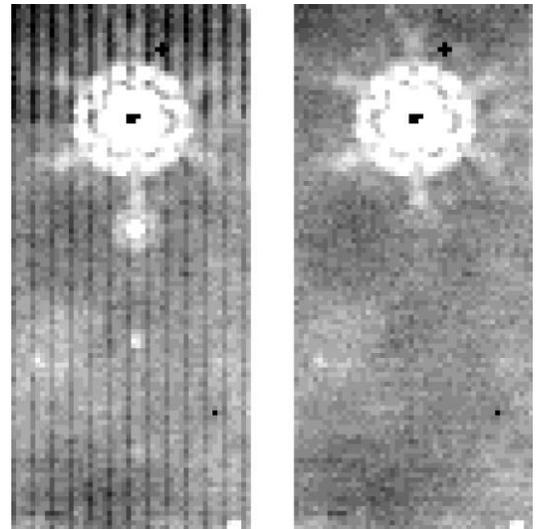


FIG. 6.—Example of bright latents and the results of the correction by applying the model. This is an ~ 10 Jy source; latents are visible for the first three intervals (left). For the purpose of calculating the response, the pixels in the source core have been replaced with the values of a model PRF projected into the array and scaled to the source flux.

known PRF to the measured source flux, and projecting the core into each BCD array at the appropriate location.

Figure 6 shows an example of the bright latent modeling and subtraction for an ~ 10 Jy source, for which the PRF core substitution has been applied.

Saturated extended sources remain a problem as we do not know the profile of a source in the saturated region. In these cases we replace the saturated array pixels, for the purpose of latent prediction, with a fixed level (4000 MJy sr^{-1}) that is moderately above the hard-saturation threshold. This provides a partial correction to the bright latents, but leaves residual latent effects, particularly for very bright compact extended objects. These have not been addressed in the current processing.

4.3. Washboard

In many AORs, the BCDs suffer from a horizontal band of depressed levels at the bottom of the array. The band varies from AOR to AOR but is usually 5 to 10 rows high and is typically $\sim 1 \text{ MJy sr}^{-1}$ in depth. Figure 7 displays this phenomenon. In a given AOR the depressed band is usually nearly constant over the AOR. Uncorrected, the band causes a pronounced rippling “washboard” pattern in the resulting mosaics due to the $1/6$ array shift (~ 21 rows) between BCDs in a scan leg.

The effect is not well understood or characterized, but it appears not to be a reduction in responsivity as the depth does not appear to increase with background brightness in a given AOR (in fact there is some evidence the effect decreases with increas-

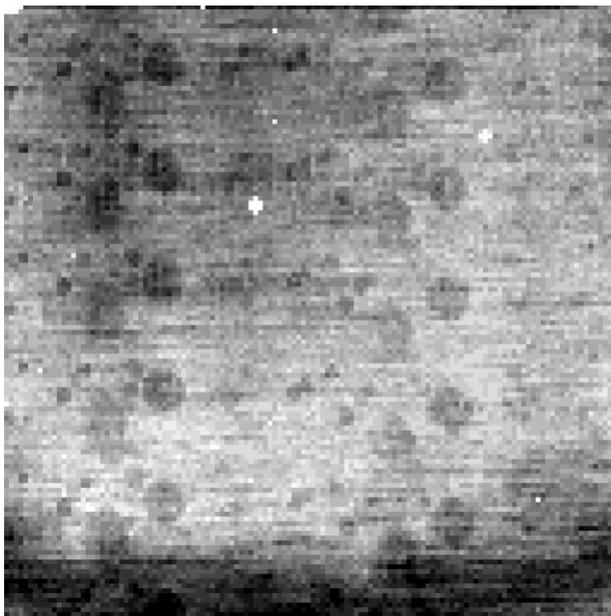


FIG. 7.— Quiescent image of an entire AOR as described in § 4.4. The washboard artifact is the depressed band at the bottom of the image. Numerous dark latents are also visible. The large faint depression at upper left is also commonly seen in the quiescent images for MIPS GAL AORs, but leaves no apparent artifact in the mosaics.

ing background). The effect seems to be more pronounced in AORs with lower mean background levels; it is mostly absent in AORs nearest the Galactic center ($|l| \lesssim 10$) and is nearly ubiquitous at large longitudes ($|l| > 50$).

We treat the washboard artifact as a constant offset artifact over the course of a single AOR. There is no dedicated correction applied, but the mitigation is accomplished as part of the dark-latent correction procedure described below.

4.4. Dark Latents

Dark latents are localized, long-lasting depressions in the array response created by passage across very high intensity levels. The resulting artifacts appear in the mosaics as sequences of dark spots typically $\sim 1 \text{ MJy sr}^{-1}$ in depth. The threshold level to trigger the effect is approximately equivalent to the peak brightness of a 20 Jy point source ($\sim 15,000 \text{ MJy sr}^{-1}$). Note that this is well into hard saturation. The decay rate has not been characterized, but dark latents in the MIPS GAL data are typically observed to persist from their onset to the end of the AOR, and are also occasionally present throughout an entire AOR following a prior passage across some bright source. Annealing the MIPS $24 \mu\text{m}$ array removes the effect, and for this reason anneals were performed every 12 hr in the MIPS GAL data taking to minimize buildup, but the dark latents remain a significant artifact.

The dark-latent effect is caused by both point and extended sources. The solid angle extent of the effect depends on the brightness profile of the source; for point sources, the dark latents are roughly circular regions from a couple of array pixels in diameter at the threshold triggering flux up to ~ 12 pixels for the brightest point sources. Dark latents due to extended sources can of course be any size, and there are examples tens of pixels in diameter.

Above the triggering threshold, there is no apparent correlation of the depth of the effect with incident brightness—brighter point sources produce larger, but not proportionally deeper, dark-latent effects, although a weak correlation has not been ruled out.

In the MIPS GAL data, the scanning is done with a $1/6$ -array shift between BCDs, and so a bright source will cause dark latents at five or six locations in the array in a single scan leg. Subsequently, and for the same reason, any sky location visited by one of the latent regions on the array will generally also be visited by the others, so in the resulting mosaics the dark-latent effects are reinforced in the averaging, producing long chains of dark spots in the images. Figure 8 shows an example of such a sequence of dark latents. There are sufficient bright sources in the MIPS GAL coverage area that virtually every AOR is affected by dark latents.

The depth of the observed depressions ranges from ~ 0.5 to 2 MJy sr^{-1} , but these values are measured at low-background brightness levels. There is some indication that the depth correlates linearly with the local background levels, and the relation suggests that the dark-latent effect is a reduction in responsivity

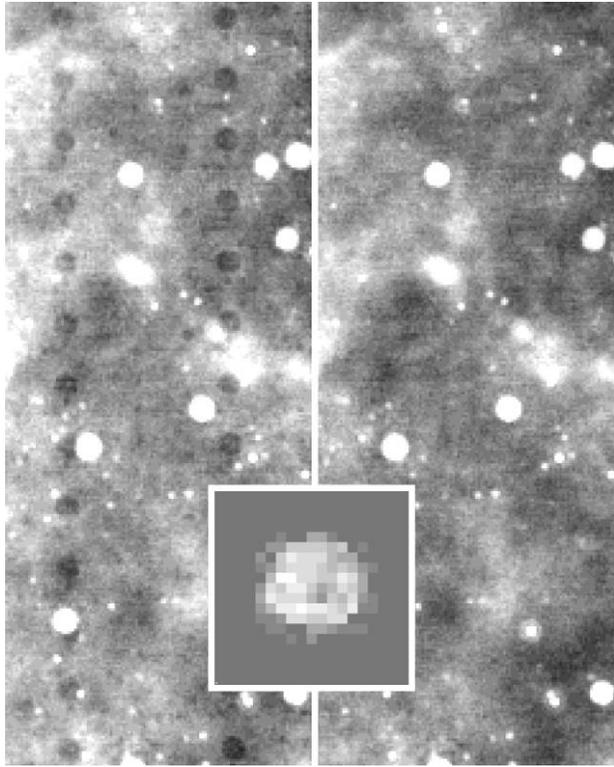


FIG. 8.—*Left*: Extracted section of an uncorrected MIPS GAL mosaic showing a chain of prominent dark latents. The latents were caused by a 200 Jy source, which crossed the array in two successive scan legs, causing latents at 10 locations on the array. *Right*: The same section after the postage-stamp quiescent image correction. *Inset*: Example of a correction image. There are 10 such images for the prominent latents, one for each of the locations on the array.

by $\sim 1\%$. This result is inconclusive, however, so to be conservative we treat the correction as an offset rather than a scaling factor. As the corrections are derived from the lowest background levels, the dark latents may be undercorrected in regions of very bright background.

The persistence of the dark latents means that the BCDs within a given AOR can be contaminated either from passage across a bright source at some time prior to the AOR, or they can be affected by observation of a bright source during the AOR. If the bright source was observed prior to the AOR, the effect is present as a quasi-constant artifact throughout the AOR. If the source is observed within the AOR, the effect persists from the observation to the end of the AOR. In practice, the two cases, dubbed “preexisting” and “newly-occurring” dark latents, are treated as similar but separate corrections.

The basic approach is to exploit the long duration of the latent effects, characterizing them by identifying persistent deviations from the local background. Using the median image over the affected BCDs is a common method for this, but the Galactic plane has far too much bright extended structure to obtain a useful result. Instead, we make use of the fact that the AORs do spend a significant, though minority, portion of

the time covering fairly low-background, quiescent regions of the sky (the endpoints of the scans for MIPS GAL I AORs are at Galactic latitudes $+1$ and -1 , approximately). In place of the median we use, in effect, an average over a range of N th-percentile values for each pixel, where N ranges from ~ 2 to 10.

For a given section of the array containing a latent, a stack is made of the BCDs affected by the latent. For each array pixel in the section, the values over the stack are extracted and sorted. The lowest seven values are omitted as assumed outliers, and a set of the next highest values are averaged together. This process should exclude any localized, structured emission the pixel observes, and preserve any systematic deviation from a local true sky mean. For stacks containing the entire AOR, we average 25 values (out of ~ 1000 BCDs in a typical AOR), and for smaller subsets of BCDs we average 10% of the total. Calculating this low-end average value for each of the pixels in the array section creates the “quiescent image” for that latent. Figure 7 shows an example quiescent image for an entire AOR.

The newly-occurring dark latents are corrected first. The latent-causing point sources are identified from the lists assembled for the bright latent correction, for which we assume every source above the threshold (~ 20 Jy) will create dark latents, and the extended sources are identified from examining the mosaics of the individual AORs, preliminary MIPS GAL mosaics, or the BCDs themselves. For each latent-causing source, and each location where it impacts the array, a postage-stamp region of the array encompassing the latent is selected, and the quiescent image is calculated from the downstream BCDs. A planar function is fit to the image perimeter, and subtracted to give a zero-point correction at the boundary. The image is then reversed and applied to the downstream BCDs as an additive correction; to avoid border effects, the correction is applied only for pixels that were saturated in the original source, as the dark latents are only triggered by brightness levels far into saturation. Figure 8 displays an example of the correction for a newly-occurring latent.

Following this correction, the only long-duration artifacts in the BCDs will be any preexisting dark latents and the washboard artifact, if it is present. To correct these effects, the quiescent image is calculated over the entire array for the whole AOR, and the image median is then subtracted as the zero-point correction. In this case, we do not know in advance where the dark latents fall on the array, so the correction image is applied over the whole array to all of the BCDs in the AOR.

A third type of persistent artifact reported in the MIPS 24 μm data, an extremely long-duration bright latency, should also be corrected with this procedure, although this effect has not been specifically seen in the MIPS GAL data set.

4.5. Masked Artifacts

Several types of artifacts are not correctable in practice and are simply edited out of the BCD images where they occur. Figure 9 displays a few examples of these artifacts.

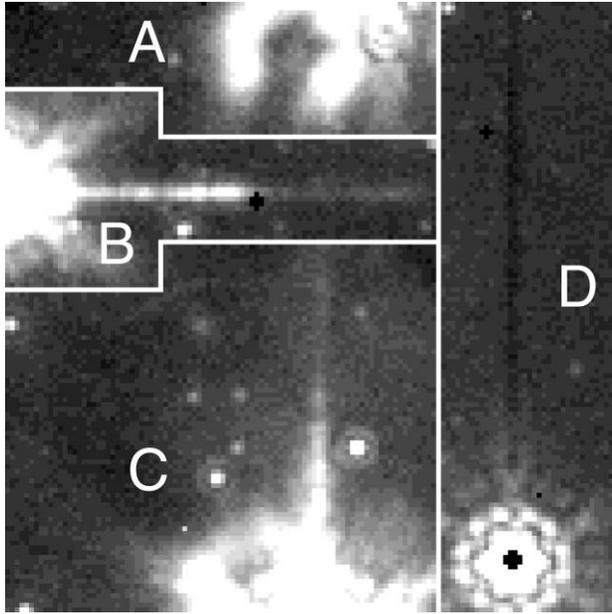


FIG. 9.—Composite BCD array showing several miscellaneous artifacts in their original locations on the array. *A*: Loop or horseshoe artifact due to source off the top edge. *B*: Horizontal streak artifact due to bright source just off the left side of the array. *C*: Both loop and vertical streak from a single source off the bottom edge. *D*: Depressed columns due to presence of a bright point source on the array.

4.5.1. Edge Streak

When a bright point source falls near the edge of the array, an internal reflection causes an illuminated streak to fall on the array adjacent to the source. If a source falls near the sides, the streak covers adjacent rows, and if it falls near the top or bottom edge, the streak covers the nearest columns. The extent of the streak depends on the brightness of the source: fainter sources will cause a streak that covers about a quarter of the array width or less, and for very bright sources the streak can extend across the BCD. The threshold source flux for producing a visible artifact is less than 0.5 Jy, so there are a very large number of these artifacts in the MIPS GAL data set.

Sources that fall near the array sides produce streaks that appear in a succession of BCDs owing to the scan motion, which is roughly parallel to the sides of the array. These streaks will generally be retained in the mosaics, as they appear in enough samples to suppress outlier rejection. These artifacts are masked out by a semiautomated procedure. The sources are identified from the mosaics, in which the general appearance of the streak is as a one-sided “diffraction spike” extending cross scan from the source. A catalog of these sources is created in which a severity value is assigned for each source that represents the extent of the streak. The masked region is specified as 1/4, 1/2, or the full width of the array, depending on the severity value. The BCDs for a given AOR are stepped through, and the near-edge sources are mapped into each BCD array. For those that fall in a zone equivalent to 2–10 pixels off the array

edge, the appropriate region is omitted by flagging the region in the corresponding BCD mask.

For sources near the top or bottom edge of the array, the streak only appears in a single BCD, and the outlier rejection usually eliminates most of the streak from the mosaics. In about 10% of the occurrences, however, the streak leaves a visible residual. In these cases, the offending BCD is identified, and the affected region is masked interactively.

4.5.2. Edge Loop

Similar to the edge streak artifact, when a source brighter than ~ 10 Jy falls near but just off the edge of the array, a large loop- or horseshoe-shaped artifact appears in the array adjacent to the source. It is assumed that this is also due to some internal reflection. The shape and orientation of the loop depend strongly on the position of the source relative to the array. The consequence is that the outlier rejection appears to remove most of the loop artifact as it is not a constant effect on the sky for a given scan passing near the source. The rejection is not complete, however, and the typical appearance of the loop residual in the mosaics is as a ragged patch of brighter background near the source. These artifacts, numbering a few hundred over the survey, are also identified in the mosaics and masked interactively.

4.5.3. Row and Column Corruption

In addition to the jailbar effect, a bright source (~ 10 Jy or higher) on the array produces another artifact in which the rows or columns containing the brightest parts of the PRF are occasionally elevated or depressed. Typically, the pixels in the columns containing the source are depressed, whereas the rows may be either depressed or elevated. The conditions that generate these effects have not been determined and so no automated masking procedure is practical. These artifacts are also identified in the mosaics and masked interactively.

This effect is also seen in *Spitzer* IRAC Si:As data (Reach et al. 2006).

4.5.4. Asteroids

When asteroids are encountered, they are usually observed only in a single scan pass, due to proper or parallactic motion between the scan passes covering a given sky position. Even when they are covered in two scans in a single AOR, they typically move enough between scans to be recorded in slightly separate positions. As there are generally two passes over each line of sight on the sky in the MIPS GAL observations, an asteroid encounter usually occupies half the data at that sky location. Although the outlier rejection is designed to retain all the data in such a situation, the algorithm seems to eliminate the bright central portion of the asteroid images (possibly due to Poisson noise skewing the brightness levels in the asteroid-containing data), leaving a doughnut-shaped feature in the mosaics. Figure 10 shows an example of the appearance of an asteroid.

The peculiar appearance of the asteroids in the mosaics allows them to be identified interactively and removed. Where the doughnut shape is seen, the component BCDs are examined to confirm the presence of a transient source. A catalog is made of the asteroid observations, which lists the AOR, sky coordinates, and effective angular diameter for each asteroid. In the correction sequence, an asteroid masking step is included, wherein the BCDs for each AOR are stepped through, the relevant asteroid coordinates are projected into the arrays, and are masked out of the BCDs where they appear.

Approximately 850 asteroids were identified and masked from the data. Nearly all (95%) were found between $l = 350^\circ$ and $l = 20^\circ$, or within about 15° of the ecliptic plane, which crosses the Galactic plane at $l \approx 6^\circ$.

It is presumed that the partial rejection that allows the asteroids to be identified will not apply below some flux threshold, and so faint asteroids will likely not be flagged and removed in this process. Also, it is possible that the angular motion of some asteroids is small enough that both scans covering a sky location will observe an asteroid at nearly the same coordinates, and so it will appear as an ordinary point source. The asteroid removal is therefore not assumed to be complete.

5. OVERLAP CORRECTION

The overlap correction is a procedure to bring all the BCDs in the MIPS GAL survey to a mutually consistent overall level by adding an optimized scalar offset to each BCD. There are

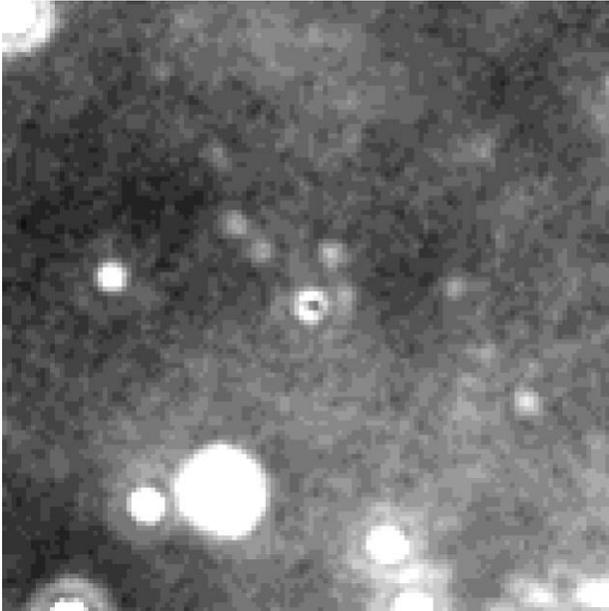


FIG. 10.—Example of the appearance of an asteroid, at center, in the MIPS GAL mosaics before masking is performed. The asteroid is present in only one of two scan passes at this location, and the outlier rejection in the mosaic generation removes the bright central portion of the asteroid image. This asteroid has a flux of ~ 30 mJy.

generally no large-scale AOR-to-AOR differences in background levels (due, for example, to Zodiacal differences or inaccuracies in the Zodiacal model). However, there occur localized deviations in the levels of individual BCDs from their neighbors or among small groups of BCDs. For example, BCDs containing very bright point sources typically suffer from an elevated overall mean brightness (usually residual droop error), and BCDs near the beginnings or ends of scan legs often do not match the levels of neighboring legs. Also, as a consequence of the jailbar effect and its correction, many BCDs have a level discontinuity across the triggering source, as described previously; the levels of these BCDs are corrected with the overlap procedure, albeit with special handling.

5.1. Basic Overlap Algorithm

The correction method is to apply a scalar offset adjustment to each BCD, determined by a least-squares minimization of the sum of the squares of the net level differences between overlapping BCDs. Let d_{ij} be defined as the original difference in the overlap region between BCD i and BCD j (which we define as the difference in the overlap-region medians). Let Δ_i and Δ_j represent the applied offsets for the respective BCDs. The net difference is then $d_{ij} + \Delta_i - \Delta_j$. The sum of the squared differences is

$$\chi^2 \equiv \sum_{ij} (d_{ij} + \Delta_i - \Delta_j)^2, \quad (7)$$

where the summation is over only the i, j pairs that actually overlap.

Minimizing with respect to each Δ , we have, for each BCD i ,

$$\frac{\partial \chi^2}{\partial \Delta_i} = 2 \sum_j (d_{ij} + \Delta_i - \Delta_j) = 0 \quad (8)$$

or

$$\sum_j (\Delta_i - \Delta_j) = - \sum_j d_{ij}. \quad (9)$$

With N_i defined as the total number of overlaps for BCD i , we have

$$N_i \Delta_i - \sum_j \Delta_j = - \sum_j d_{ij} \quad (10)$$

or, in matrix form,

$$\mathbf{A} \cdot \Delta = \mathbf{B}, \quad (11)$$

where

$$A_{ij} \equiv \begin{cases} N_i, & \text{if } i = j \\ -\delta_{ij}, & \text{if } i \neq j \end{cases}, \quad (12)$$

$\delta_{ij} \equiv 1$ if BCDs i and j overlap, zero otherwise, and

$$B_i \equiv -\sum_j d_{ij}. \quad (13)$$

This simple system of equations is ill conditioned as the solution is not unique: any global offset to the set of corrections Δ will also satisfy equation (9) for all BCDs i . This is not a conceptual difficulty as we can apply some reasonable condition on Δ as a constraint (perhaps minimizing the sum of the squares of the Δ s), although solving this set of equations for a very large number of BCDs may be computationally problematic. A more serious problem is that the least-squares condition for a single BCD (eq. [9] or [10]) is also individually satisfied with any arbitrary offset added to the local corrections. This results in the overall solution being very poorly constrained against large, spurious gradients in the resulting background levels: if the solution to the equations requires extreme global variations in the Δ s, then that is what we get, although it is unlikely to be physical in any ordinary circumstances. We typically see this exaggerated “ramping” result when the BCDs suffer from some systematic inhomogeneity, such as the washboard effect.

5.2. Damped Overlap

The MIPS 24 μm pipeline processing removes instrumental bias, and so the background levels of the MIPS GAL pipeline-produced BCDs are, collectively, a measure of the absolute background. For the overlap correction, therefore, we want to preserve as much as possible the global background as-is, while removing BCD-to-BCD differences. To mitigate the potential ramping problem, the approach we take is to include the offsets themselves in the minimized function, along with the overlap differences, with the intent of suppressing large magnitudes of Δ . We redefine χ^2 as

$$\chi^2 \equiv \sum_{ij} (d_{ij} + \Delta_i - \Delta_j)^2 + \sum_k \beta_k (\Delta_k)^2, \quad (14)$$

where k is summed over all BCDs, and the β_k are selected to control the strength of the damping effect.

The least-squares minimization gives, for each BCD i ,

$$\sum_j (d_{ij} + \Delta_i - \Delta_j) + \beta_i \Delta_i = 0. \quad (15)$$

If a set of Δ s local to BCD i satisfies the original overlap condition given by equation (8), then this equation is nearly satisfied also, provided β_i is sufficiently small and the net local Δ offset is also small (i.e., no ramping in the solution). Conversely, a set of Δ s that satisfies this relation also nearly satisfies equation (8). That is, a solution to this damped set of equations

is also very nearly a solution to the basic overlap correction conditions, provided the β s are small enough. The asymmetry of Δ_i in equation (15) assures that the overall level of the local Δ cannot be set arbitrarily as part of a global solution because a uniform offset to the Δ on the left hand side no longer cancels. Also, note that if the BCDs satisfy the minimization condition (eq. [9]) without correction, i.e., $\sum_j d_{ij} = 0$, then equation (15) is also satisfied with zero correction, and so this procedure will not add deviations to already-matching data.

In essence, this is an overlap solution that allows small residual level differences while largely preserving the global background levels. The fractional deviation of any Δ_i from the “true” value is approximately the fractional change in the coefficient of Δ_i on the left hand side of equation (15) from that of equation (8), or β_i/N_i (provided $\beta_i \ll N_i$). In other words, the residual overlap difference for BCD i and overlapping BCD j after the correction will typically be on the order of $(\beta_i/N_i)d_{ij}$.

In practice, β_i is set such that the damping effect is independent of the mean number of overlaps. We set $\beta_i = \alpha N_i$, where α is a single free parameter specifying the strength of the damping. The damped solution is thus simply a change to the diagonal elements of the overlap correction matrix, from $A_{ii} = N_i$ to $A_{ii} = N_i(1 + \alpha)$. The optimal value of α depends on the acceptable fractional residual overlap difference. For MIPS GAL, typical uncorrected overlap differences are, at most, on the order of 2 MJy sr^{-1} , whereas the smallest difference visually detectable is about 0.1 MJy sr^{-1} , so α for MIPS GAL should be no more than about 0.05. Figure 11 shows a section of a MIPS GAL mosaic containing a “bad” overlap difference, and the overlap correction results for several values of α . For $\alpha = 1.0$, we see very little correction, as expected; for $\alpha = 0.1$, the overlap differences have been largely eliminated.

For MIPS GAL, we used $\alpha = 0.04$. Figure 12 shows a before-and-after example of the damped overlap correction for the same mosaic as Figure 11 but over a larger region.

One benefit of the damping procedure is that the resulting matrix is well conditioned, and for the MIPS GAL data, it is very sparse, so computationally it is a straightforward solution even for very large numbers of BCDs. In the current release, the first quadrant, fourth quadrant, and Galactic center regions were processed separately; each contained from 150,000 to 190,000 BCDs.

5.3. Outliers

Typical overlap differences are 1–2 MJy sr^{-1} , but numerous fairly isolated BCDs deviate from their neighbors by up to tens of MJy sr^{-1} , usually from a residual droop when crossing very bright point sources, or residual jailbar level errors. The damped overlap correction itself cannot correct these cases effectively, as the residual overlap difference would typically remain a visible artifact; worse, the solution would tend to drag neighboring BCDs up or down to match the deviant BCD.

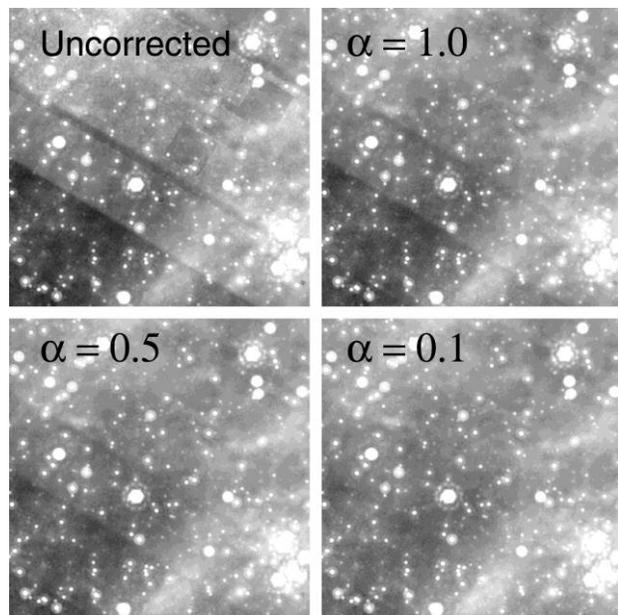


FIG. 11.—Typical bad overlap difference in a MIPS GAL mosaic, and the effects of different damping parameters on the overlap correction solution.

The remedy is to turn off the damping effect for these BCDs, which amounts to setting the diagonal elements of the matrix back to N_i , so that the levels of those BCDs are allowed to float to match the neighboring BCDs, and also to decouple overlapping BCDs so that their resulting offsets are not influenced by the outlier. That is, for all overlapping BCDs j set $A_{ji} = 0$ (while keeping $A_{ij} = -1$ to tie the outlier to its neighbors), and make corresponding changes to N_j and B_j . If there are only a small percentage of outlying BCDs this approach does not affect the overall overlap solution.

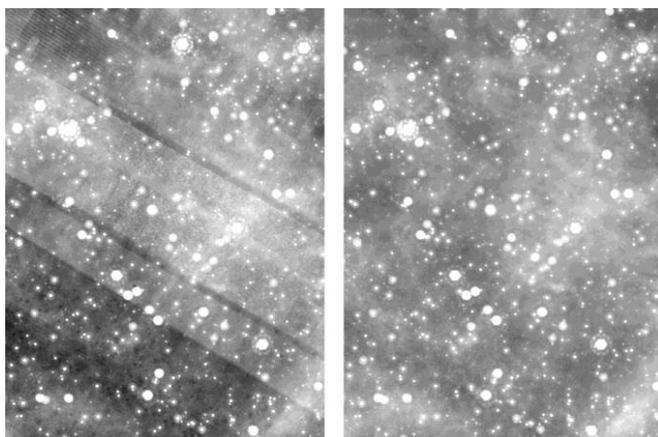


FIG. 12.—*Left*: Section of a MIPS GAL mosaic before corrections are applied to the BCDs. Overlap errors, jailbars, and numerous dark latents are visible. *Right*: Same section after the BCD corrections are applied.

5.4. Overlap Correction and the Jailbar Discontinuity

As previously described, the correction of the jailbar artifact often leaves a level discrepancy above and below the triggering source. These level errors are corrected with an adaptation of the overlap correction procedure.

First, a catalog of point sources and compact bright extended objects that cause such discontinuities is compiled by examining either a preliminary set of mosaics or the AOR mosaics. The sources that cause a discontinuity are typically very bright, >20 Jy, but some much fainter sources also can cause problems. Consequently, no automatic process can account for the problem, and an interactive effort is needed to identify them.

The catalog of triggering sources is an input to the overlap correction routine. For BCDs on which a source that causes the jailbar discontinuity falls, the array is divided into two parts, above and below the source. Each segment is processed as an independent image in the correction routine. The overlap offsets are calculated, and for the partitioned BCDs, the results are applied in the appropriate regions.

6. RESIDUAL ERRORS IN THE ARTIFACT CORRECTIONS

As Figures 11 and 12 demonstrate, the corrections described here effectively remove evident visual artifacts from the mosaics. The issue remains, however, to what extent these procedures may leave residual level errors, or perhaps introduce them, at a level below visual detectability, and what effects such errors may have on fluxes measured from the images.

All the corrections save for the bright latents are additive adjustments, and therefore possible residual flux errors are primarily due to uncertainties in the local background levels. These, in turn, can be estimated from the detectability limits of residual errors in the corrections.

For extended objects, the relevant corrections are those that affect regions larger than a point source, namely the jailbar, washboard, and overlap corrections, and the bright latent correction, which is most evident for point sources but is applied to all data. The jailbar correction is applied to entire BCDs, but is periodic on a length scale smaller than a point source, so residual errors should not affect extended-object photometry, and in any case the overall level is determined by the overlap correction. The overlap correction is a scalar applied to each BCD (or segment of a BCD) so any resultant errors will be present as hard-edged discontinuities, and this will be true also for errors in the washboard correction. Empirically we find that the lower limit of visual detectability for extended, isolated-edge discontinuities is about 0.1 MJy sr^{-1} in quiescent regions and about 0.2 MJy sr^{-1} in regions of bright structured emission. The BCDs are spaced about an arcminute apart along a scan leg, so the maximum spurious local gradient that can be introduced is from 0.1 to $0.2 \text{ MJy sr}^{-1} \text{ arcminute}^{-1}$, depending on

the local backgrounds, and this gives an upper limit to the flux error of an extended object as a function of its angular extent.

The bright latent effect is about a percent of the incident brightness, so for extended objects there should not be a flux determination issue even without a correction. To be more quantitative, however, we can estimate the error of the bright latent model by noting that we can eliminate, from visual detectability, latent effects of point sources up to at least about 1 Jy. For a source of this flux, the peak brightness is about 700 MJy sr^{-1} , and the first interval peak latent brightness is about 7 MJy sr^{-1} . The minimum detectability of a pointlike source is $\sim 0.5 \text{ MJy sr}^{-1}$, so the correction is accurate to about 10% or better. The residual error is then at most about 0.1% of the incident brightness. For an extended object, the residual latent error due to its own latent effects, or effects from neighboring objects of equivalent brightness, is therefore negligible. For compact extended objects in close proximity ($1'$ or so) to a much brighter object, possible residual latent errors may occasionally be significant.

For point sources, the possible residual errors due to the jailbar, washboard, and overlap corrections are those due to uncertainty in the background as these corrections are applied to BCD-sized regions and not to localized areas. The background uncertainty in a region used for source fitting, typically less than $1'$ in diameter for faint sources, is roughly equal to the overall deviation due to any steplike discontinuities included, so we expect at most a $0.1\text{--}0.2 \text{ MJy sr}^{-1}$ error in background levels. This is negligible for all but the very faintest detectable point sources.

The bright and dark-latent corrections do apply localized corrections so we need to consider the possible effects on point source peak brightnesses directly. For the bright-latent corrections, the only significant problem is the case where point sources (or compact extended objects) happen to coincide with the locations of bright latent residuals of much brighter nearby sources. As the total solid angle of bright-source latents is a very tiny fraction of the survey area, statistically the problem should be negligible. For specific sources, of course, the local neighborhood needs to be examined for these sources of error.

The dark latents can be from a few to tens of pixels in diameter and corrections for them are ubiquitous in the images. An isolated dark-latent correction can be in error by up to about 0.5 MJy sr^{-1} without visual evidence but the chain pattern of the dark latents makes the detectability limit $0.2\text{--}0.3 \text{ MJy sr}^{-1}$. Undetectable dark-latent correction errors are therefore potentially problematic only for the very faintest point sources and extended objects. The solid angle covered by dark-latent corrections is larger than that for bright-source latents, but still a small fraction of the total, so the problem should not be statistically significant.

7. DISCUSSION AND SUMMARY

The objective of producing a completely automated pipeline for artifact correction proved elusive; particularly due to the

need for generating comprehensive source lists for the latency corrections, and the need for visual inspection of the resulting mosaics to identify and mask stray light artifacts and remedy other residual problems.

Nevertheless, all artifacts common in the MIPS 24 μm data have been addressed and corrected in the MIPS GAL data products.

Jailbars. The mitigation for the jailbars is the closest we have to a “black box” correction procedure. The one common situation for which jailbars are left uncorrected is where large extended saturated regions cover the BCD, in which case there is insufficient information to characterize the readout medians. These are masked out of the data.

Bright latents. The basic latent modeling effectively removes latent images if the incident brightness is represented accurately. For unsaturated sources, there is rarely any residual from the latent subtraction visible in the mosaics. For saturated sources, however, the process of estimating the core profile with a projected PRF has many ways to result in poor subtraction.

First, the modeled PRF profile must match the actual source profile. The PRF model we use was created from a composite of point source images located across the array. The actual PRF, however, apparently changes in width by up to 5%–10% depending on position in the array. In the mosaics, many bright sources have a prominent latent-subtraction residual that is both symmetric and is distributed evenly between peaks and valleys. This indicates an accurate position and flux, but an inexact PRF shape.

Second, the catalog position of the source must be very precise. A discrepancy of half an arcsecond leaves visible residuals. Further, the coordinates must be correct relative to the actual position of the source on the array, and not necessarily the true position on the sky. BCD pointing error is typically a half arcsecond, and this can result in residuals even if the absolute source position is accurate. The approach we take is to optimize the source positions and fluxes by applying the latent-subtraction mechanism and minimizing the residuals from the first latent. This works well in quiescent backgrounds, but if there is structured emission overlapping the first latent, the optimization usually fails. We also use an interactive procedure to optimize the parameters by hand when the automated method fails, but this requires the problem sources to be identified visually.

Finally, compact saturated extended sources are always poorly corrected, as we do not know the profile of the source. Very compact sources are a particular problem as they can masquerade as point sources; if they are corrected as point sources, the subtraction residual is a trail of dark spots because the PRF is distinctly narrower than the source, and if not, the source leaves a trail of bright latents because the (minimal) default correction for saturated pixels routinely underestimates the peak brightness.

Dark latents. On the whole, dark latents are well removed from the images. The primary remaining problem is a failure

to recognize dark-latent-causing sources. A number of chains of small or faint dark latents remain in the mosaics for this reason. Point sources are usually reliably flagged, but a few at the threshold latent-causing brightness slip through. All latent-causing extended objects need to be selected by visual inspection so this is another source of omission.

Another problem, though rarer, is the situation where dark-latent-causing sources are dense enough that a new dark latent in an AOR coincides with the occurrence of other dark latents in the AOR. The correction image for a dark latent can easily be corrupted by the presence of other dark latents at the same location on the array. There are a few cases of dark-latent residuals in the mosaics where this seems to be happening.

The issue remains of whether the dark latents should be corrected as scaling factors rather than offsets. We agree that there is convincing evidence that this is the case but we opt for caution before applying a multiplicative correction over Galactic hill

and dale. This will be investigated for inclusion in later releases of the MIPS GAL images.

Masked artifacts. Masking is a regrettable solution for artifacts, because it can leave distinctly poorer coverage; the edge streak and loop artifacts in particular can necessitate removing half the data coverage where they occur. The deletion of coverage near bright sources is the most visible residue of the presence of artifacts in the data; very noisy patches, unevenness, and even data holes are common near bright sources as they often cause artifacts to some level in every BCD that crosses them.

This work is based on observations made with the *Spitzer Space Telescope*, which is operated by the Jet Propulsion Laboratory, California Institute of Technology under a contract with NASA. Support for this work was provided by NASA through an award issued by JPL/Caltech.

REFERENCES

- Engelbracht, C. W., et al. 2007, *PASP*, 119, 994
 Gordon, K. D., et al. 2005, *PASP*, 117, 503
 Kelsall, T., et al. 1998, *ApJ*, 508, 44
 Reach, W. T., et al. 2006, *IRAC Data Handbook Version 3.0* (Pasadena: SSC), <http://ssc.spitzer.caltech.edu/irac/dh>
 Rieke, G. H., et al. 2004, *ApJS*, 154, 25
 Werner, M., et al. 2004, *ApJS*, 154, 1