Publications of the Astronomical Society of the Pacific, 134:094101 (17pp), 2022 September https://doi.org/10.1088/1538-3873/ac8a18

SMC-Last Mosaic Images

D. R. Mizuno[1,4], Kathleen E. Kraemer[1], T. A. Kuchar[1], and G. C. Sloan[2,3]
[1] Institute for Scientific Research, Boston College, 140 Commonwealth Avenue, Chestnut Hill, MA 02467, USA; mizunod@bc.edu
[2] Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA
[3] Department of Physics and Astronomy, University of North Carolina, Chapel Hill, NC 27599-3255, USA

Received 2022 February 24; accepted 2022 August 16; published 2022 September 13

Abstract

We present mosaic images of the Small Magellanic Cloud (SMC) observed with the Spitzer IRAC 3.6 μm and 4.5 μm bands over two epochs, 2017 August 25–2017 September 13 and 2017 November 24–2018 February 12. The survey region comprises ~30 deg^2 covering the SMC and the Bridge to the Large Magellanic Cloud. The region is covered by 52 ~ $1\overset{\circ}{.}1 \times 1\overset{\circ}{.}1$ tiles, with each tile including images in each band for both separate and combined epochs. The mosaics are made in individual tangent projections in J2000 coordinates. The angular pixel size is $0\overset{''}{.}6$ with a resolution (FWHM) of ~$2\overset{''}{.}0$. We describe processing to correct or mitigate residual artifacts and remove background discontinuities. The mosaic images are publicly available at the Infrared Science Archive.

Unified Astronomy Thesaurus concepts: Small Magellanic Cloud (1468); Infrared astronomy (786)

1. Introduction

The Small Magellanic Cloud (SMC) is a nearby metal-poor dwarf galaxy. Its nominal distance is 62 kpc (e.g., de Grijs & Bono 2015; Graczyk et al. 2020), but its structure is complex, with a foreground component at a distance of ~55 kpc from the Sun on the east side, likely due to tidal interactions with the Large Magellanic Cloud (LMC) (Nidever et al. 2013; Subramanian et al. 2017) and another component several kiloparsecs behind the southwest regions of the galaxy (e.g., Scowcroft et al. 2016; Yanchulova Merica-Jones et al. 2021). The metallicity of the SMC depends on the sample observed. Cepheids give [Fe/H] = −0.68 (Luck et al. 1998), while for the red-giant population, the mean [Fe/H] = −0.97 (Choudhury et al. 2020). Rubele et al. (2018) find that [Fe/H] ranges from ~−0.6 for the youngest population to ~−1.6 for the oldest.

Thus, the SMC serves, on the one hand, as a laboratory for studying the evolution of a population of stars as a whole, at a known distance and at a metallicity significantly lower than the Galaxy. On the other hand, the SMC is a complex and dynamic metal-poor galaxy close enough to our own that it can be studied in great detail.

Understanding the chemical evolution of this galaxy requires observations of the stars both as they die, when they enrich the SMC with freshly fused elements and dust, and as they form. Because stars evolving to and from the main sequence are variables, a proper study of those populations requires multi-epoch surveys. The Optical Gravitational Lensing Experiment (OGLE; e.g., Udalski et al. 1997, 2008) has monitored the SMC at a near-daily cadence from 1997. The OGLE-III catalogs include data through 2008 and have identified over 19,000 long-period variables in the SMC, including over 2,000 variables likely on the asymptotic giant branch (Soszyński et al. 2011).

Photometry in the near-infrared is more sensitive to the cooler objects evolving to and from the main sequence. The VISTA Magellanic Cloud Survey (VMC) surveyed 45 deg^2 in the SMC at Y, J, and K_s, with 12 epochs at K_s, starting in 2011 (Cioni et al. 2011). The SMC region of the VMC covers the Bar, which includes most of the well-known H II regions in the galaxy, and the Wing, which extends to the east, toward the LMC. An additional 20 deg^2 covers part of the Bridge between the Magellanic Clouds. The Infrared Survey Facility mapped the central square degree at J, H, and K_s (Ita et al. 2018). While not as deep as the VMC, it has better temporal coverage, with over 100 epochs from 2001 to 2017.

The mid-infrared is best suited to study the most embedded objects, the stars that are actually forming or dying. Using the Spitzer Space Telescope, Bolatto et al. (2007) mapped a region covering the Bar and part of the Wing of the SMC in the Spitzer Survey of the Small Magellanic Cloud (S^{3}MC). The survey covered roughly 10 deg^2 using all of the filters available on the Infrared Array Camera (IRAC; 3.6, 4.5, 5.8, and 8.0 μm; Fazio et al. 2004) and on the Multiband Imaging Photometry for Spitzer (MIPS; 24, 70, and 160 μm; Rieke et al. 2004). The IRAC observations were obtained in 2005 May. S^{3}MC was

---

[4] Author to whom any correspondence should be addressed.

expanded on by the Surveying the Agents of Galactic Evolution program (SAGE; Meixner et al. 2006), which had initially observed the LMC. SAGE-SMC observed a 30 deg^2 area covering the entire Bar and Wing of the SMC and extending east to the Bridge with all seven IRAC and MIPS filters (Gordon et al. 2011). The temporal coverage included two epochs, with the IRAC data in 2008 June and September. After the end of the cryogenic portion of the Spitzer mission, the SAGE-VAR survey obtained four additional epochs from 2010 August to 2011 June at 3.6 and 4.5 μm in 3 deg^2 centered on the Bar of the SMC (Riebel et al. 2015).

Our Spitzer program (Program ID number 13096) built on this legacy by adding two more epochs at 3.6 and 4.5 μm covering the entire 30 deg^2 footprint of the SAGE-SMC survey in late 2017 and early 2018. This program, entitled Spitzer's Last Look at the Small Magellanic Cloud, or SMC-Last for short, when combined with the previous surveys, provides a minimum of four epochs at 3.6 and 4.5 μm creating a temporal baseline of over 9 yr covering the entire SMC and its surroundings. Up to nine epochs are available in the core of the SMC spanning a period of over 12 yr.

Section 2 describes the two new epochs of observations with Spitzer of the SMC that form the basis of SMC-LAST. Section 3 presents an overview of the resulting image products. The following sections cover some of the steps taken when processing the data, with corrections for background levels in Section 4, corrections for artifacts in Section 5, masking for artifacts in Section 6, and other issues in Section 7. Section 8 provides some final thoughts.

2. Observations

We mapped ∼30 deg^2 of the SMC, including the Bar, the Wing, and the Bridge that extends toward the LMC, at 3.6 and 4.5 μm (Program ID 13096) using Spitzer's Infrared Array Camera (IRAC; Fazio et al. 2004). The Astronomical Observation Requests (AORs) were based on those used for the 2008 maps for SAGE-SMC from Gordon et al. (2011). The observations were taken in the high-dynamic range mode (HDR), with 0.4 and 10.4 s integrations, which optimized the sensitivity to both bright and faint objects.

Each of the 29 primary AORs consisted of a 28 × 14 raster with 146.″4 and 292.″8 steps (120 and 240 pixels) for the rows and columns, respectively, covering regions that are ∼1.°1 × 1.°1. The 29 raster scans needed for one iteration of the SMC map were intended to be made with the same roll angle, and the two iterations of the map were to be executed three months apart. That would provide a 90° roll, and the requested observation dates were chosen to avoid gaps within the maps. This strategy gives a factor of at least 2 in coverage in a single map, enabling the removal of random effects like cosmic rays and bad pixels, as well as systematics such as latents from saturations or scattered light (e.g., Hora et al. 2004, 2006, 2008). Because many of these artifacts preferentially affect either rows or columns in the array, the 90° roll from the second epoch allowed us to better mitigate systematic effects such as column pull-down effects from bright sources, as well as providing an additional factor of 2 in coverage.

For the first map (epoch 1), the observations were made largely as planned, taking place over the course of 18 days, 2017 August 25–September 13. A small rotation is present in the AORs in the Bridge relative to those nearer the core, from the rotation of the IRAC field of regard between execution of the initial and final AORs. Due to data-volume constraints and limited availability of the ground station, the observations for the second epoch spanned 78 days, starting 2017 November 24 and ending 2018 February 12. The resulting change in roll angle over this period left gaps among the planned AORs. Working with the Spitzer Science Center enabled us to add several smaller AORs using Director's Discretionary Time to fill in many of those gaps. Figure 1 illustrates the final IRAC coverage for the two epochs.

3. Overview of Image Products

The SMC-Last mosaics comprise, for each epoch and channel, 52 FITS image files, each covering a region approximately 1.°06 × 1.°06, with an angular pixel size of 0.″6. The images are made in tangent projections in J2000 coordinates, with the projection center at the center of each image and a rotation angle of 0°. The images are aligned in rows 1° apart in decl., and in each row the images are approximately 1° apart in R.A., with the overall alignment of the images chosen to best cover the survey region. The image-to-image overlaps are 3.′5 in decl. and a minimum of 1.′6 in R.A. Figure 2 shows a schematic of the image boundaries superposed on a composite image of 4.5 μm data for the combined epochs.

The data comprise image sets for each of the two channels in both epochs, plus each channel in a combination of both epochs. In addition to the image data, each image file has corresponding FITS files containing coverage maps (the number of data pixels influencing each mosaic pixel) and uncertainty maps.

Table 1 provides a summary of the properties of the image products.

The resolution of the mosaics has been investigated by creating composite point-source images for each plate and fitting a 2D Gaussian function to the composite images. The mosaic sources have been selected by mapping sources from the 2MASS "6x" catalog (Cutri et al. 2012) (and the 2MASS catalog (Skrutskie et al. 2006) where the 6x catalog lacked coverage, longward of about R.A. 01$^{\mathrm{h}}$52$^{\mathrm{m}}$), using sources in the K-magnitude range of 12.0–17.0. The matching mosaic point sources are aligned with the centroid pixel and averaged together. The 2D Gauss fitting yields median FWHM values of ∼2.″0 over all the data, in both x and y, corresponding closely

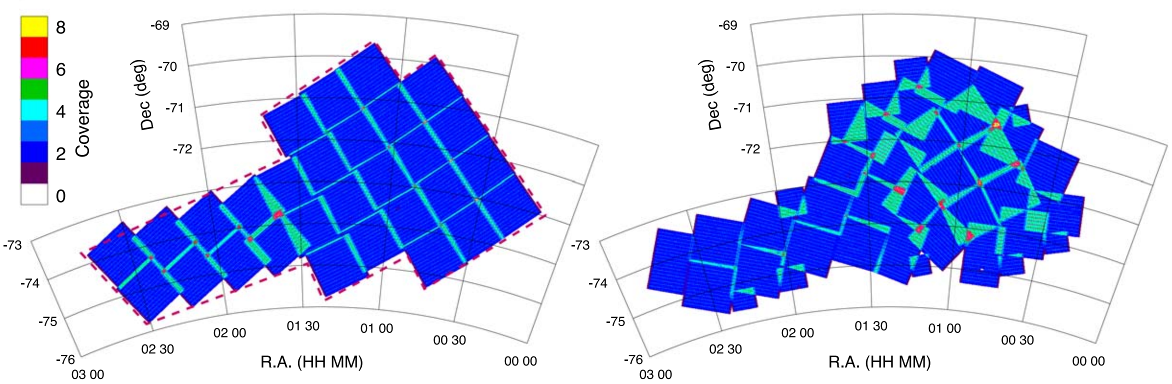

Figure 1. Final coverage of the SMC-Last program. (Left) Epoch 1, 2017 August 25–September 13; (right) epoch 2, 2017 November 24–2018 February 12. The SAGE-SMC survey region is shown as the dark red dashed outline. “Coverage” is the number of 10.4 s integrations observing each location in the region.

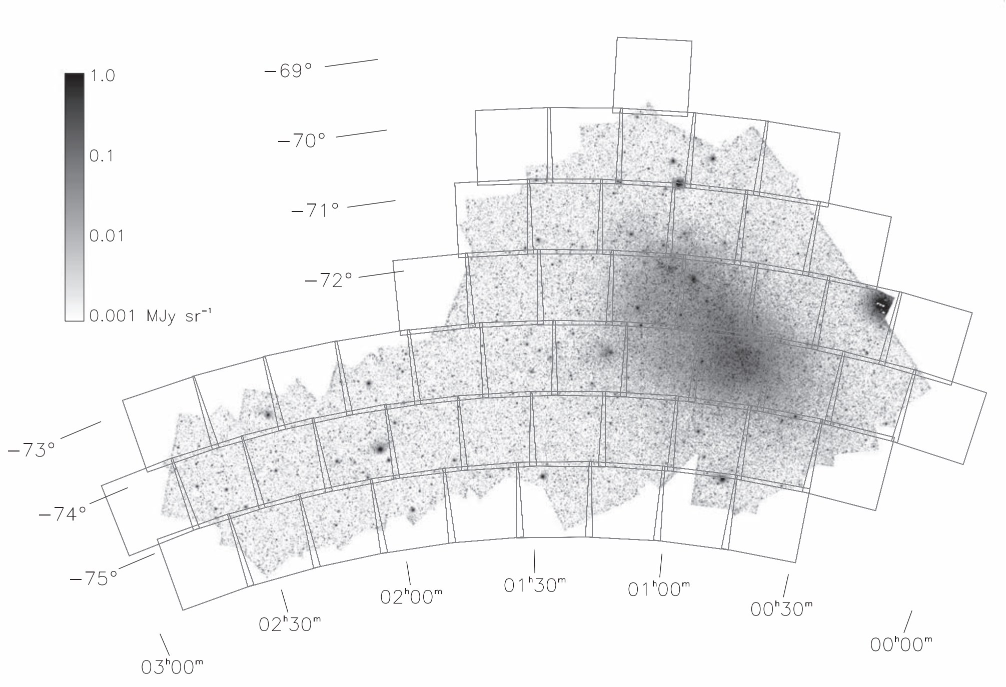

Figure 2. The 52 FITS image boundaries superposed on a composite of the combined-epoch data for channel 2.

to the R.A. and decl. axes, with a mosaic-to-mosaic rms of about $0\farcs1$. The centroid pixel alignment does smooth the point-source profile somewhat; we estimate a pixel-centered profile would have an FWHM $\sim 0\farcs05$ smaller.

The mosaics have been constructed using the mosaic utility in the Mopex software package provided by the Spitzer Science Center (Makovoz & Khan 2005). The Corrected Basic Calibrated Data (CBCD) pipeline products provided by the Spitzer Science Center have been used throughout as the input data source to construct the mosaics. The supplied CBCD data products have been corrected for detector linearization and flat-fielding, calibrated to flux units (MJy sr^{-1}), have had dark current removed and the post-BCD pointing refinement applied (although see the Appendix for details of the pointing refinement used for our data), and have been flagged for radiation hits, saturation effects, latent images, stray light, and other artifacts (IRAC Instrument & Instrument Support Teams 2021) A zodiacal component has been removed with the dark frames, but not from the data, so the CBCD products have the difference included. The zodiacal model values supplied with the CBCD products indicate that the zodiacal levels typically vary by less than 0.005 MJy sr^{-1} over any given AOR, and this is absorbed

Table 1
SMC-Last Mosaic Properties

Property	Epoch 1	Epoch 2	Combined
Dates	2017 Aug–Sep	2017 Nov—2018 Feb	⋯
Median coverage depth	2	2	4
Pixel–pixel noise[a] I1:	0.043–0.10	0.044–0.092	0.032–0.091
(MJy sr^{-1}) I2:	0.034–0.050	0.033–0.058	0.024–0.054
Background level[a] I1:	−0.005–+0.006	−0.008–+0.005	−0.007–+0.006
(MJy sr^{-1}) I2:	−0.003–+0.002	−0.001–+0.005	−0.002–+0.004

Global Properties	
Pixel size	0.″6
Angular resolution	∼2″
Observed area	30 deg^2
No. of plates	52
Plate size	1.°06 × 1.°06
Bands	3.6 μm and 4.5 μm
Integration mode	High-dynamic Range
Integration times[b]	$t_{\rm exp} = 0.6$ and 12 s, $t_{\rm int} = 0.4$ and 10.4 s

Notes.

[a] Pixel–pixel noise: standard deviation in a 51 × 51 pixel box (30 arcsec^2) in a source-free region. Background level: mean level in the same box. Estimated from one region in the SMC Bridge and one just outside the Core.

[b] In HDR mode, the two nominal exposure times of 0.6 and 12 s lead to integration times of 0.4 and 10.4 s.

into the background-level adjustments described below. The 10.4 s integrations only have been used in the mosaics, except where, as noted below (Section 7.1), the 0.4 s integration data are used to fill in saturated pixels for very bright sources.

The options for coadding in the mosaics include weighting by the uncertainty values, which according to the Mopex handbook (Makovoz & Khan 2005) is not recommended, weighting by integration time, or equal weighting. We have used equal weighting, but as we have used the 10.4 s integrations almost exclusively, this is effectively weighting by integration time. The exception is the saturation substitutions, discussed in Section 7.1.

The CBCD products have been processed further prior to the construction of the mosaics as described in the remainder of this paper. In the following, the term "pipeline" refers to the CBCD processing pipeline and "BCD" refers to a single frame of CBCD data for the 256 × 256 detector arrays. Channels 1 and 2 are the 3.6 and 4.5 μm data, respectively.

4. Background-level Adjustments

Figure 3 shows a composite image of the epoch 1, channel 1 data using the CBCD products without further processing. The scaling is linear to emphasize the low-level backgrounds. It is apparent that there is a significant nonastronomical variation in the background levels from AOR to AOR, with an overall scatter of about 0.1 MJy sr^{-1}. This variation indicates that this data set does not yield absolute background levels or variations at a length scale of the typical AOR coverage region (about a degree). For the purpose of point-source extraction from the mosaics, we would like to remove the discontinuities at the AOR boundaries.

If these level variations can be characterized by scalar offsets at the BCD level (as opposed to a true gradient being added at either the BCD or AOR level), then in principle the backgrounds can be brought to a common, self-consistent value with an overlap-matching procedure in which each BCD is given a scalar offset, with the offset values simultaneously optimized with a least-squares algorithm to minimize each BCD's overlap differences with its neighbors. In practice, however, an overlap correction procedure that exactly minimizes the overlap differences tends to produce large "ramping" excursions in background levels because such a routine responds to small systematic variations in BCD backgrounds, and there is no constraint on the magnitude of the corrections.

Also apparent in Figure 3 is that the AOR backgrounds not only have a generally scalar differential offset, but typically also an overall gradient, either increasing or decreasing in level over the duration of the AOR, plus in many cases a steep increase or decrease in background levels at the start of the AOR. The approach we take is to "detrend" the backgrounds for each AOR separately (described in Section 4.2) and then apply a modified version of the BCD overlap-matching procedure to resolve residual BCD-to-BCD level differences (Section 4.3). To mitigate potential ramping in the overlap-matching procedure, processing is first applied to reduce systematic pixel response variations and artifacts that persist for the duration of an AOR. This is described in the following section.

4.1. Median-image Subtraction

The median BCD array images for the AORs (i.e., the median value of each array pixel over the duration of an AOR) typically show artifacts that are persistent for a large fraction of the AOR. Figure 4 shows a few examples. Also of concern, for the overlap-matching procedure, are possible small systematic variations across the array. To mitigate these effects, a form of the median image is subtracted from each BCD: The AOR is divided into two sections chronologically (i.e., for a 1 hr AOR, each section consists of 30 minutes of data). The median BCD image is calculated for each section, and then the overall level of each is normalized by subtracting the median value of the image, giving a mean value close to zero. A correction image is formed by assigning each pixel as the smaller magnitude of the two median images. This correction image is visually compared to the first and last BCDs in the AOR; if the artifact in the correction image represents a transient and is not present in both BCDs, the process is repeated with the AOR divided into

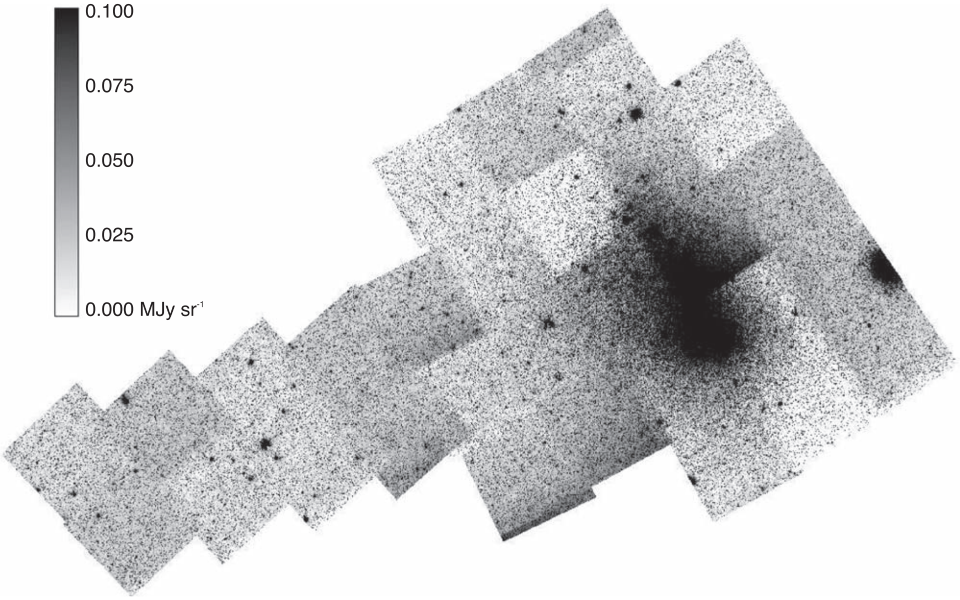

Figure 3. Composite of epoch 1, channel 1 raw CBCD data, without the background-level adjustments described in Section 4.

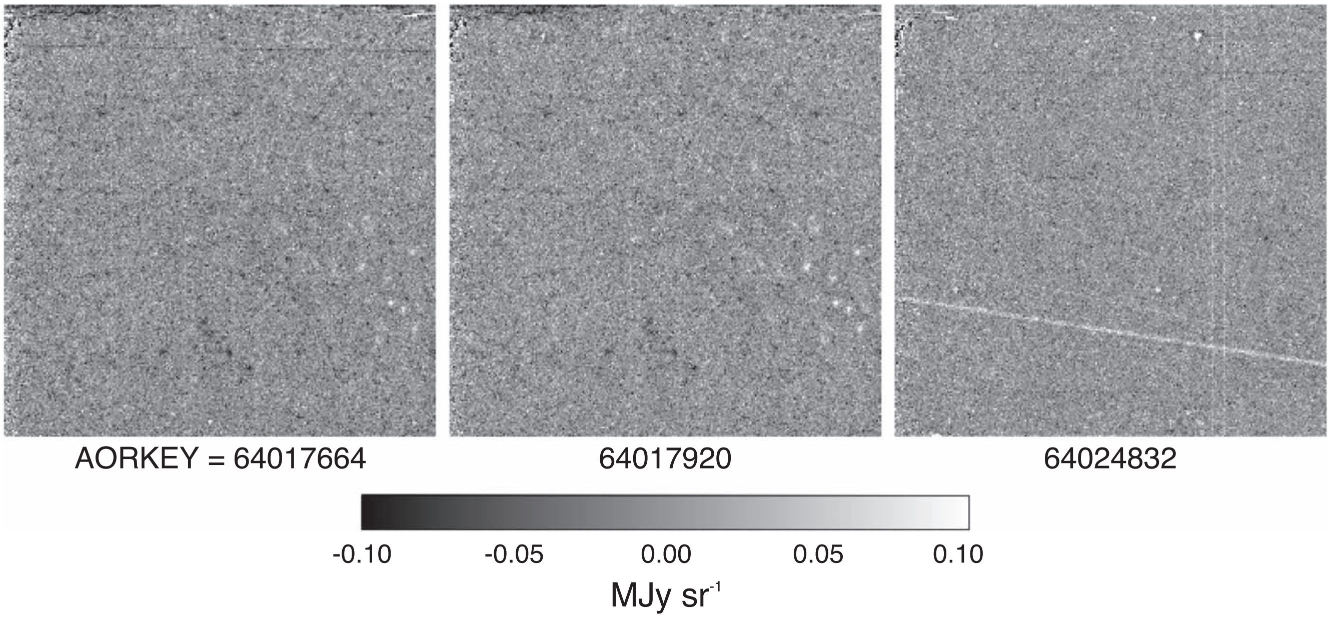

Figure 4. Examples of median images for three AORs, all for channel 1. Typical are the spots seen at lower right in the first two examples. These may be latent images of bright point sources but they are not flagged by the CBCD pipeline. The streak in the example on the right is likely the result of the array slewing across a very bright source.

three or four sections as necessary to eliminate the transient artifact from the correction.

Two concerns arise with this procedure. The first is that the correction images have an apparent residual noise level, and so, the subtraction could increase the uncertainty in the BCD data. We have examined four AORs from each epoch in the Bridge region (where the backgrounds are fairly flat) to investigate this question, comprising $\sim$3000 BCDs in each band. For each pair of BCDs (original CBCD and median-image subtracted), most sources, plus artifacts and bad pixels, are first deleted by eliminating pixels more than 0.25 MJy sr^{-1} from the median in either BCD and then eliminating pixels within one pixel of

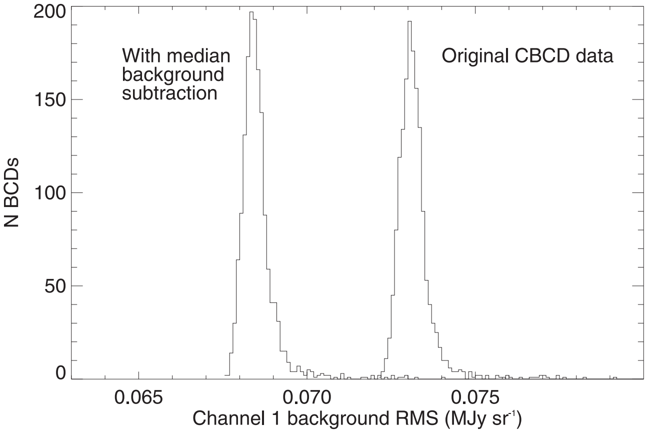

Figure 5. Histograms of background noise levels of ~3000 BCDs from four AORs in the Bridge, for the original CBCD data, and the results of the median background subtraction. This is channel 1 data from the first epoch.

those pixels. The rms of the remainder is then calculated. For the 3.6 μm data, the original CBCD data have a median rms of 0.073 MJy sr^{-1}, and the median-image-subtracted data have a median rms of 0.069 MJy sr^{-1}. For the 4.5 μm data, the results are 0.059 MJy sr^{-1} and 0.055 MJy sr^{-1}, respectively. Figure 5 shows a histogram of the results for the 3.6 μm data in the first epoch. We find thus that the rms noise levels in the BCDs actually decrease slightly, by about 5%–7%.

To investigate whether this decrease could be due to deletion of artifacts rather than an actual reduction in the background noise, Figure 6 shows median images constructed from the first half of the BCDs from AORKEY = 64024064 (channel 1) and from the second half. The third panel shows the difference between the two. If the rms reduction were due to the removal of artifacts, and the background "noise" in the median images were true noise, the difference image would show a noise level increased by about $\sqrt{2}$. On the contrary, the rms of the difference image is reduced by more than a factor of 2 (~0.033 MJy sr^{-1} in the median images and ~0.014 MJy sr^{-1} in the difference image). This indicates that the "noise" in the median correction images is primarily some true systematic pixel-to-pixel variation consistent throughout the AOR.

The second concern is that this correction will tend to remove true sky gradients from each BCD if such a gradient is uniformly present across the region covered by the AOR. While subtraction of an overall sky gradient will not occur in this step (because each BCD in an AOR is given the same normalized background correction), BCD-level flattening could potentially produce a "staircase" background artifact, although we see no evidence of this in the mosaics.

4.2. Detrending

In this procedure, the scalar BCD background levels for each AOR are adjusted separately to remove overall trending in the levels. The median of each BCD in the AOR is taken, and the results are initially fitted with a function consisting of an exponential plus a linear term:

$$y(x) = ae^{\frac{-x}{b}} + cx + d, \tag{1}$$

where the ordinate values x are the scaled sequence number of the BCD in the AOR and the coefficients are determined with a nonlinear least-squares algorithm, using uniform weighting for all the data points. Then, for values of x greater than 3 times the "time constant" b (to exclude the exponential portion), the data values above the fitted curve are given zero weight, and the values below are weighted by the distance below the curve, with the weights normalized to the maximum distance. The curve is then refit with the adjusted weights, and then the weights are adjusted again similarly with the refit curve, giving both first and second refits.

The intent is to create an estimate of the background bias function for the AOR, where the assumption is made that the observations for any given AOR will spend a significant fraction of the time covering portions of the sky at some ambient background (away from the SMC core), which we are arbitrarily assigning zero brightness, noting as above that this data set is not expected to preserve low-level large-scale extended emission.

The intended function is thus one that skims the lower bound of the BCD-median data. The first refit usually provides the appropriate results, and in a few cases, the second refit is used. Also, in a few cases where the fitting failed, the linear portion of the curve is assigned by hand. Figure 7 shows two examples of the BCD-median function, the initial fit, and the first refit for an AOR that crosses the SMC core and one in a region with flatter backgrounds. The sinusoidal appearance of the median data for the first case is caused by the raster scanning repeatedly over the core.

For each BCD in the AOR, a scalar background bias value is calculated from the fitted curve and subtracted.

The issue for this procedure is to what extent it is affecting relative extended emission and background levels, particularly at the ~1° length scale of the AOR-sized regions. (Much smaller extended features should be unaffected by the simple linear background fitting, and we do not expect to recover background features at much larger length scales.) This will be addressed in the next section.

4.3. Overlap Level Matching

Following the detrending, all the BCDs (for a given channel) are processed with an overlap-matching procedure (Mizuno 2008) in which each BCD is given a (typically small) scalar offset to least-squares minimize the residual-level differences between overlapping BCDs to reduce or effectively eliminate background discontinuities in the mosaics. This procedure is described by Mizuno et al. (2008). This is a "damped" overlap

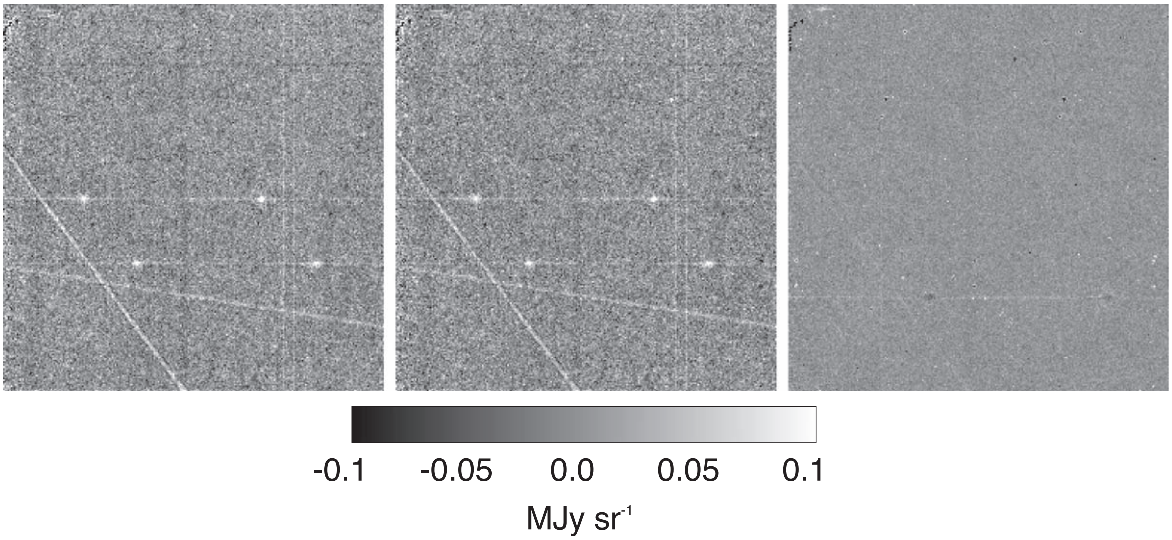

Figure 6. Median images created from the BCDs over first half (196 BCDs) of AORKEY = 64024064, channel 1 (left panel), the second half (middle panel), and the difference between the two (right panel). The rms for the median images is about 0.033 MJy sr^{-1} and for the difference image about 0.014 MJy sr^{-1}. If the apparent noise in the median images were true random noise, we would expect the rms of the difference image to be about 0.047 MJy sr^{-1}.

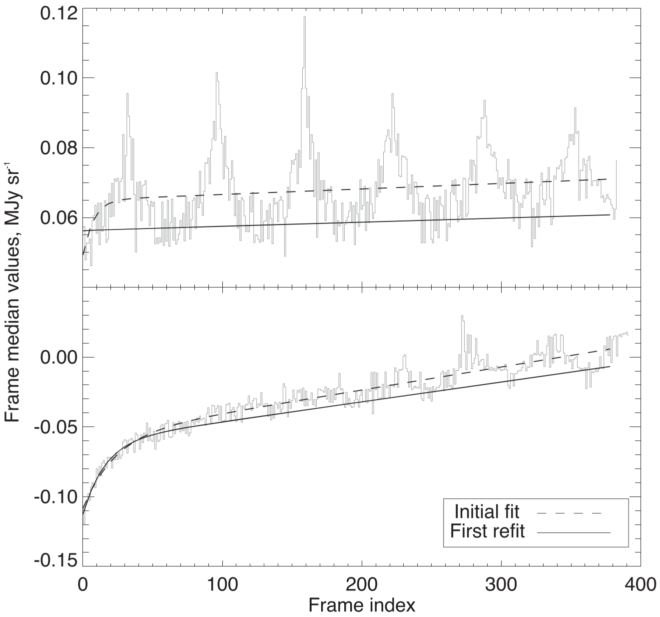

Figure 7. Two examples of the BCD-median functions, the initial fit, and first refit as described in the text. Both are channel 1. The upper case is an AOR (AORKEY = 64021760) that periodically crosses the SMC core, and the lower (AORKEY = 65213184) is in a region with flatter backgrounds.

algorithm intended to suppress the ramping effect common with an exact overlap-matching algorithm. In essence, the magnitude of the correction for a given BCD is included in the calculation of χ^2 in the least-squares algorithm, along with the overlap differences with its neighbors. The cost of the ramping suppression is, in principle, small residual-level differences in the BCD-to-BCD overlap regions, but this has not been observed in practice.

Both epochs for each channel are overlap-matched simultaneously, to give common background levels between the epochs, noting that true background levels (apart from zodiacal differences, which we are not attempting to preserve) are not going to change observably over the timescale of the two epochs.

Figure 8 shows the difference between a composite image of the unprocessed data for epoch 1, channel 1 (i.e., Figure 3) and the data processed with the detrending and overlap adjustments described above. (The median-image subtraction is omitted because it does not affect the overall background levels, but the systematic pixel value adjustments over each AOR cause visible features along the scan rasters in the difference image that obscure subtle background-level differences). Ideally we would see nothing but level variations that align with the AOR boundaries, and this is generally the case. The exception is that the backgrounds in the core of the SMC appear to be slightly oversubtracted; the arrows in the figure indicate the faint boundary visible. The magnitude of the apparent oversubtraction is around 0.01–0.02 MJy sr^{-1}. This effect is absent for the case where the detrending alone is included, so it seems to be a consequence of the overlap-matching procedure. While the detrending procedure would be expected to depress backgrounds for AORs that have higher true mean backgrounds than their neighbors, the overlap-matching procedure is not expected to have any particular overall bias, and so we cannot necessarily conclude that these results are a relative oversubtraction of the background levels in the SMC core. The alternative is to postulate that there is some systematic elevated

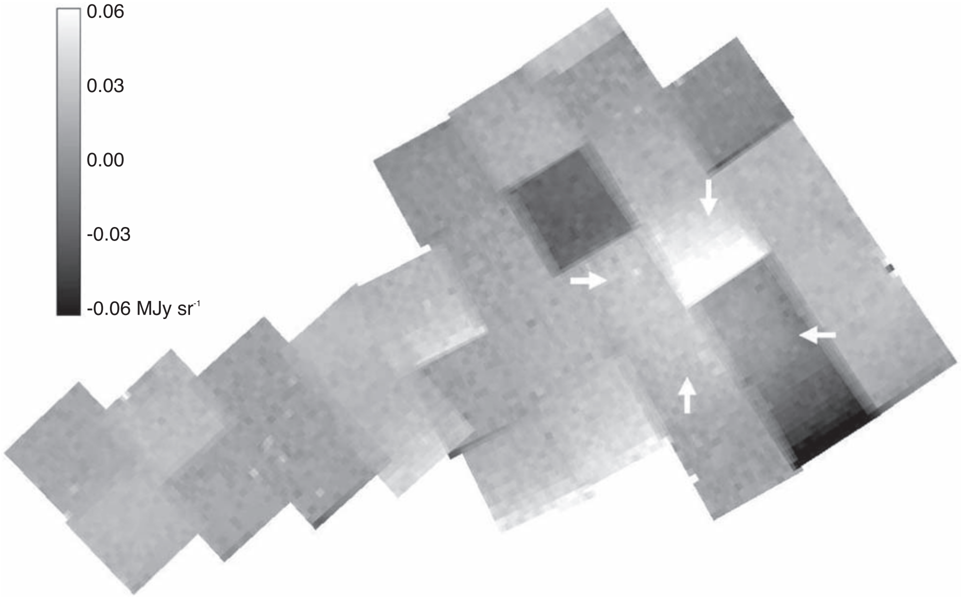

Figure 8. Difference image between the raw CBCD data (i.e., Figure 3) for epoch 1, channel 1, and the data processed with the detrending and overlap-matching procedures. The arrows show the apparent boundary of the SMC core.

bias in the raw data for the SMC core that roughly scales with the true background levels. We note that the IRAC channel 3 and 4 Si:As detectors had just such a scattered-light issue and could also be present at low levels for the InSb detectors of channels 1 and 2.

5. Artifact Corrections

5.1. Column Pull-down and Pull-up Corrections

The column pull-down artifact is a phenomenon in which entire columns that contain a very bright source are depressed in intensity. The nature of the depression varies from very uniform along the affected columns, which allows for a correction, to highly irregular, which can only be masked, and even very uniform examples typically show some irregularity at the very top and bottom of the array. The width of these artifacts generally scales with the brightness of the source, and somewhat surprisingly, the wider artifacts from very bright sources tend to be more uniform than the narrower examples.

While the CBCD pipeline includes a correction for column pull-down, examples remain either uncorrected or incompletely corrected, and we address these residual artifacts.

Column pull-up is a similar phenomenon in which entire columns have an elevated intensity. In contrast to the pull-down artifact, however, pull-up artifacts have extremely uniform elevated levels along the columns and can nearly always be corrected. For these cases, there is seldom any obvious triggering source, although many seem to be associated with a point source at the very top or bottom of the array, and others seem to be associated with the first latent images of bright point sources. The column pull-up artifacts can be from one to a few columns in width and in some cases are up to about a quarter of the array wide (these band-like artifacts are almost always associated with first latent images).

The correction for these artifacts exploits the circumstance that the backgrounds in most of the survey region are nearly flat, generally lacking any significant structured extended emission. The basic approach is to apply a scalar offset to the regions of depressed or elevated columns to match the overall median for the BCD. Specifically, the procedure is to identify the boundary columns of the depressed or elevated regions, thus dividing the BCD into groups of columns with uniform level offsets. The median of each group is given a scalar offset to match the global median for the BCD, with boundary columns adjusted individually to match the global median. Note that while narrow artifacts are corrected to the presumably nonelevated background levels over the rest of the BCD frame, for the band-like pull-up artifacts, no assumption is made about the "truth" background, and the correction just sets the column region level offsets to a common value. The subsequent background overlap-matching procedure described above adjusts the background levels to agree with the BCD's neighbors.

Very narrow pull-up artifacts are easily identified and corrected in software if the elevation is about 0.05 MJy sr^{-1} or more. The remaining pull-up and all pull-down artifacts are corrected, if possible, by hand as they are encountered from inspections of the mosaics and BCDs. The by-hand corrections also allow ad hoc masking of irregular portions of the pull-down artifacts. Figure 9 shows a few examples of the pull-down artifacts and the resulting corrections, and Figure 10 likewise for pull-up artifacts.

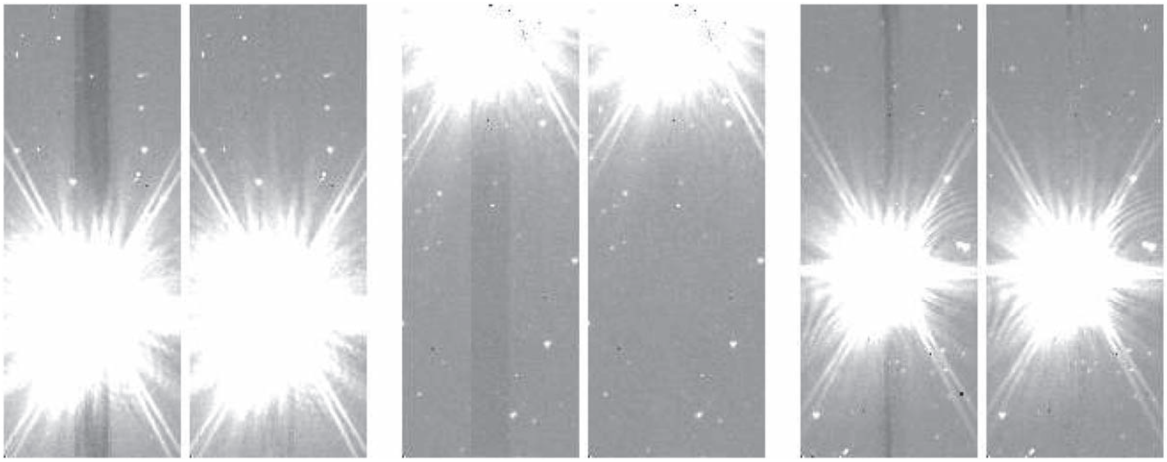

Figure 9. Examples of the pull-down artifact and their corrections. Note that the correction for the case on the right leaves a residual at the top of the array; this is subsequently masked.

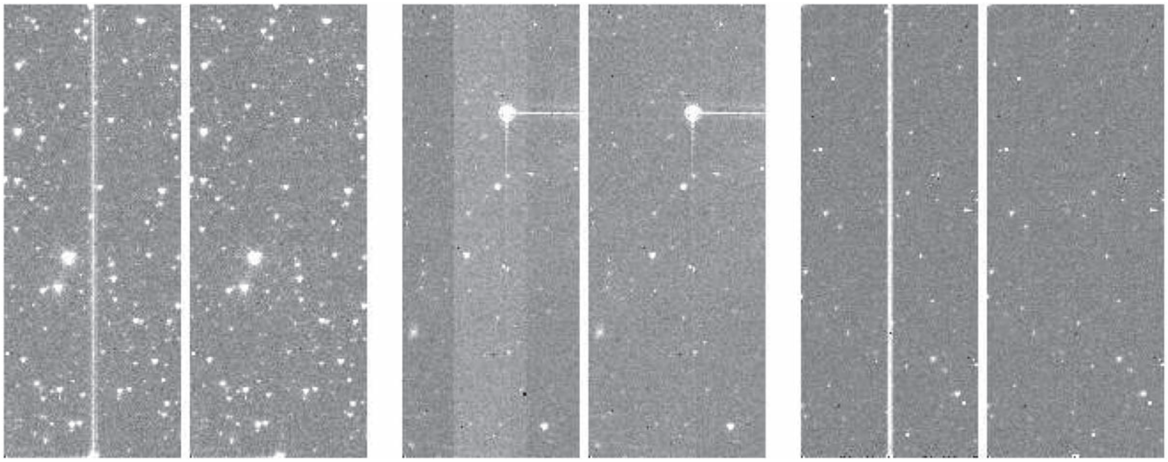

Figure 10. Examples of the pull-up artifact and their corrections. The bright source in the band-like artifact at the center is the first latent of a bright point source. Note that the left and right cases are associated with faint point sources at the bottom of the array.

5.2. Latent Corrections

In channel 1, latent images from bright point sources can persist for minutes or tens of minutes (i.e., dozens or hundreds of subsequent BCDs). For channel 2, nearly all latent images fade to undetectability within approximately a minute (in this data set). The CBCD pipeline masks the latent images, but in channel 1, for the long-duration latents, the CBCD pipeline also commonly continues the latent masking long after the latent has become undetectable.

Also, the pipeline misses many instances of latents, in both channels, particularly when they fall near the very edge of the array. The cutoff appears to be about 10 pixels from the edge. For the 3.6 μm data, over all the AORs in the survey, the pipeline has flagged 634 latent sequences further than 10 pixels from the array edge; we would therefore expect about 112 sequences ($\sim$15% of the total) nearer the edge, whereas the pipeline has flagged 20. For the 4.5 μm data, similarly, the pipeline flagged 1190 sequences further than 10 pixels from the edge, and only 24 nearer the edge, where we would expect more than 200 (and the pipeline flagged none within 10 pixels of the top and right edges).

The long-duration channel 1 latents typically decrease in brightness rapidly and nonlinearly for the first few BCDs and then fade nearly linearly until they become undetectable. This linearity is exploited to apply a correction: for a given latent image occurrence in the linear portion of the sequence, the latent itself is estimated as the median value of the affected region of the array over a nine-BCD window centered on the given BCD, with an appropriate scalar background subtraction. The latent image estimate is then subtracted from the BCD array, and the pipeline masking for that latent image is deleted.

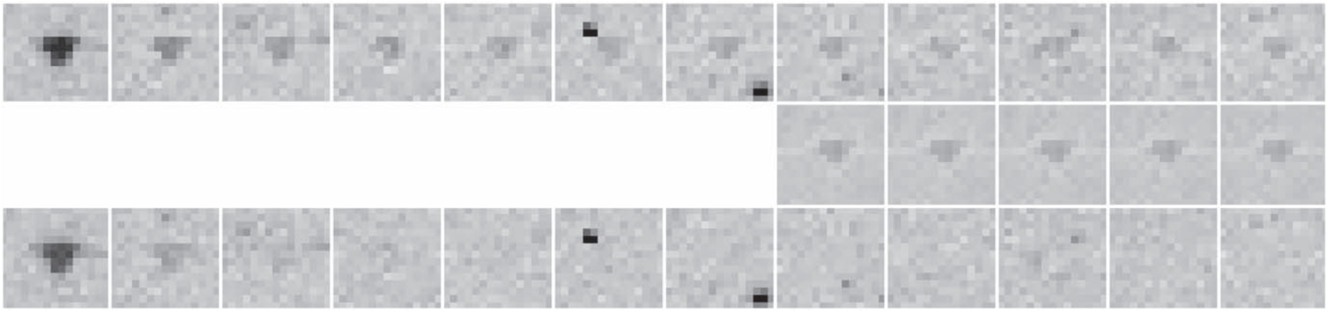

Figure 11. Example of a sequence of channel 1 latent images and the corrections. Top row: the initial latent and the 11 subsequent images at that location in the array, following the observation of a ∼3 Jy point source. In this case, the latent images persist for about 20 minutes. Middle row: the calculated latent images as described in the text. The first calculated latent is at the eighth occurrence in the sequence and is provisionally applied to the earlier occurrences. Bottom row: the affected regions after the subtraction of the latents. The first applied correction is at the fifth occurrence; the earlier latent images remain masked.

The latent correction procedure is applied with an interactive routine: For each latent sequence, regions for the correction and background estimation are selected (the pipeline tends to underestimate the size of the latents for very bright sources), then the corrections are applied and the results inspected in situ on each affected BCD, to assess the quality of the correction. Generally, the correction is effective starting at the fifth latent occurrence.

The visual inspection process also permits two additional latent correction steps: cases in which the pipeline masking has been extended long after the latent has faded, and for these the masking is simply turned off; and cases in which the pipeline misses identifying latents. For these, the short-duration cases are simply masked, and the correction procedure is applied to the long-duration cases.

While channel 2 does not have long-duration latents, the channel 2 BCDs are also inspected to identify and mask the pipeline-missed latents.

The visual inspections have found about half (64) of the "expected" edge latents for channel 1 and about a quarter (49) for channel 2 (plus, for both cases, a roughly equal number scattered over the array).

Figure 11 shows an example of a sequence of latents, the calculated latent images, and the subtracted results.

5.3. Large-scale Latents

When an extremely bright point source crosses the arrays (during slewing between frames), it can leave a streak-like latent image in the array that can persist for as long as the longest point-source latents. These cases are not flagged by the CBCD pipeline. For these latents, a procedure similar to the point-source latent corrections is applied: The boundary of the latent itself is defined, and for each latent image, the nine-BCD window median image is taken for the whole array, and, again exploiting the generally flat backgrounds in this data set, the median of this image, apart from the latent region itself, is regarded as the scalar background and subtracted. The resulting latent image is subtracted from the BCD and similarly visually inspected for the effectiveness of the correction.

6. Artifact Masking

6.1. Stray Light

The stray-light artifacts are patches of light, usually seen in a few distinctive patterns, that are confined to about the upper third of the array, presumably due to bright sources just off the upper edge of the array. The CBCD pipeline makes an attempt to predict and mask these artifacts, but it identifies only the brightest cases, and the selected masking region is typically much larger than necessary. The pipeline-flagged artifacts have been visually examined and the masking regions reset manually as necessary to accommodate the actual size and extent of the artifact.

There has otherwise been no systematic effort to identify the remaining (mostly fainter) stray-light artifacts. However, the distinctive shapes are readily apparent in the mosaics and individual BCDs and are thus masked by hand when they are located. Figure 12 shows a few examples of stray light, both caught and missed by the CBCD pipeline.

The IRAC Instrument Handbook indicates that there is also a small region off the lower edge of the channel 1 array that can cause stray-light artifacts. No cases of this have been noticed, although faint examples could have gone undetected.

6.2. Column Pull-down

While a portion of the column pull-down artifacts are correctable using the procedure described in Section 5.1, for the majority the residuals after the correction are sufficiently irregular to warrant simply masking the artifact instead. Toward that end, all point sources that are either saturated or above a peak threshold (300 MJy sr^{-1} for channel 1 and 200 MJy sr^{-1} for channel 2) are visually inspected on the BCDs for the pull-down artifact, and if present, evaluated for whether correction or masking is indicated, and then respectively applied. If masked,

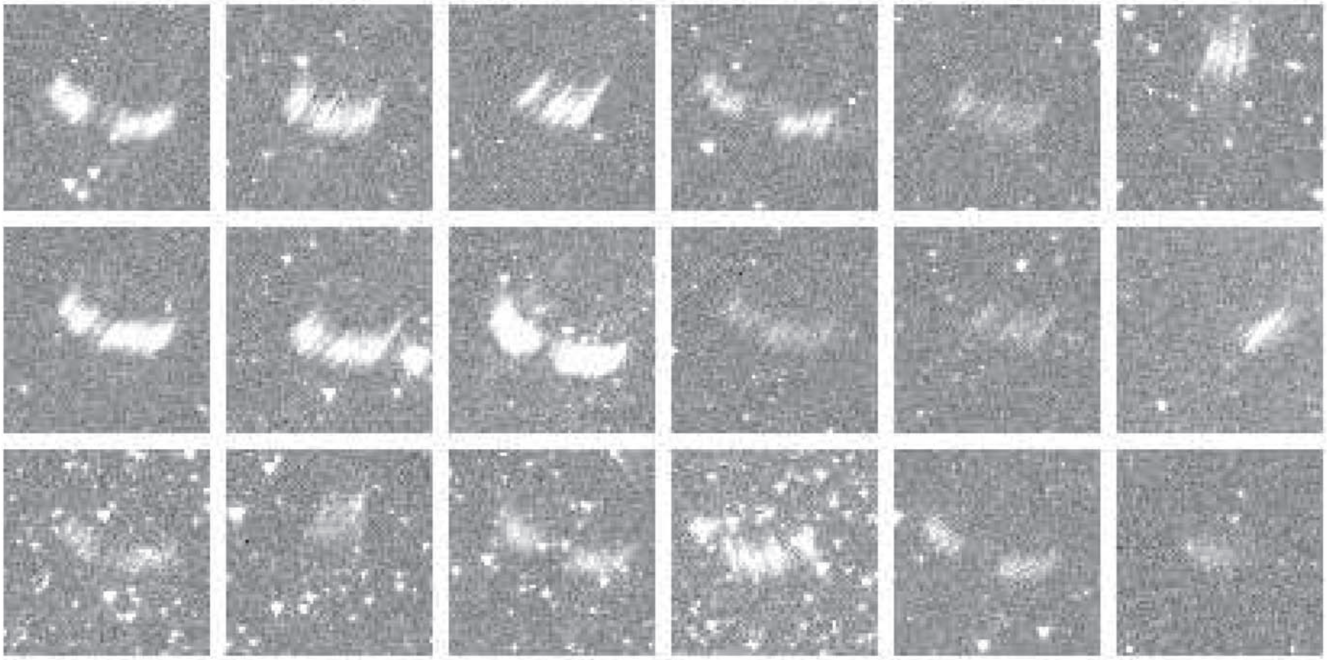

Figure 12. Some examples of the stray-light artifact. The top row shows cases caught by the CBCD pipeline, and the remainder were missed. Note that the pipeline misses some very bright cases.

the masking boundaries are selected by hand, but the triggering source itself is generally left unmasked.

A second artifact is found in cases in which, in a fraction of the columns labeled as pull-down by the CBCD pipeline, one to a few contiguous pixels have values far below the local background, from one to tens of MJy sr^{-1}. These usually occur in columns that have been flagged as pull-down and sufficiently corrected in the pipeline, as they show no residual pull-down effect. These artifacts are usually not caught in the outlier rejection mechanism in the mosaic construction, and so leave small negative “holes” in the mosaics. These artifacts are identified as pixels beyond a specified threshold level (0.5 MJy sr^{-1}) below the local background level for “pull-down” labeled columns and masked. Figure 13 shows an example of a column with this artifact and the uncorrected effect on the resulting mosaic image.

The cause of this artifact appears to be a consequence of the pipeline pull-down correction itself, as the corresponding uncorrected BCD does not show this artifact.

6.3. Charged-particle Strikes

Particle strikes can create an artifact in which a portion of the column in which the strike occurs, up to about 50 pixels above and below the strike, is corrupted such that one side is elevated and the other depressed. While the strike itself is usually either flagged and masked in the pipeline or in the outlier rejection in the mosaic creation, the corrupted pixels are not and can leave an artifact in the mosaics. These cases are masked by hand as they are identified.

Figure 14 shows a few examples of both uncorrectable pull-down artifacts and these charged-particle strikes.

7. Miscellaneous Corrections and Residual Artifacts

7.1. Saturation Substitution

The BCD pixels flagged and masked by the pipeline as bright point-source saturations have been replaced with the corresponding pixels from the short-integration (0.4s) data taken contemporaneously with the long-integration data. For this procedure, the short-integration pixel values are inserted into the long-integration BCD arrays prior to the mosaic construction, along with the substitution of the corresponding pixel values for the uncertainty data. (The pointing from the 0.4 s to the 10.4 s integrations seems to be sufficiently stable for the pixel-for-pixel substitution.) Note that we are effectively using integration time as the weighting in the mosaic construction, and so this pixel substitution will result in improper weighting for the substituted 0.4 s pixels. However, regions saturated in a given 10.4 s BCD are unlikely to have unsaturated 10.4 s data in any overlapping BCDs, and so the improper weighting is unlikely to result in improperly combined data in the mosaics.

The saturations due to true point sources are distinguished from saturations from other causes (typically charged-particle strikes on the array) by examining the data in a 4 pixel boundary around the saturation: For saturated point sources, the

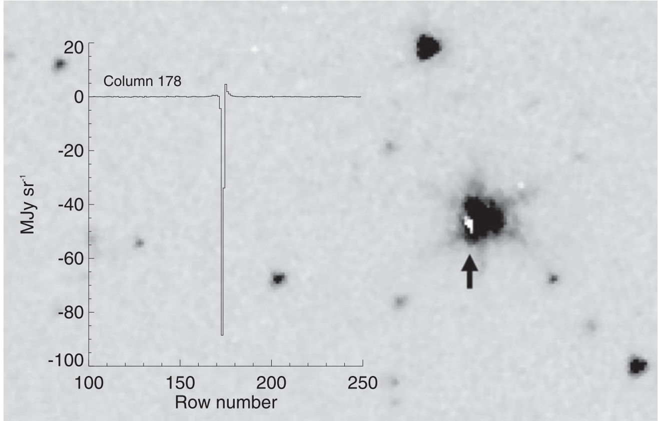

Figure 13. Example of the very negative pixel values seen in some "pull-down" columns, and their effect on resulting mosaic images. Note that this is not a data hole in the mosaic, rather a very negative value, because the outlier rejection mechanism fails to trap these cases. The affected source is at $\sim 02^{\rm h}00^{\rm m}$, $-74°52'$.

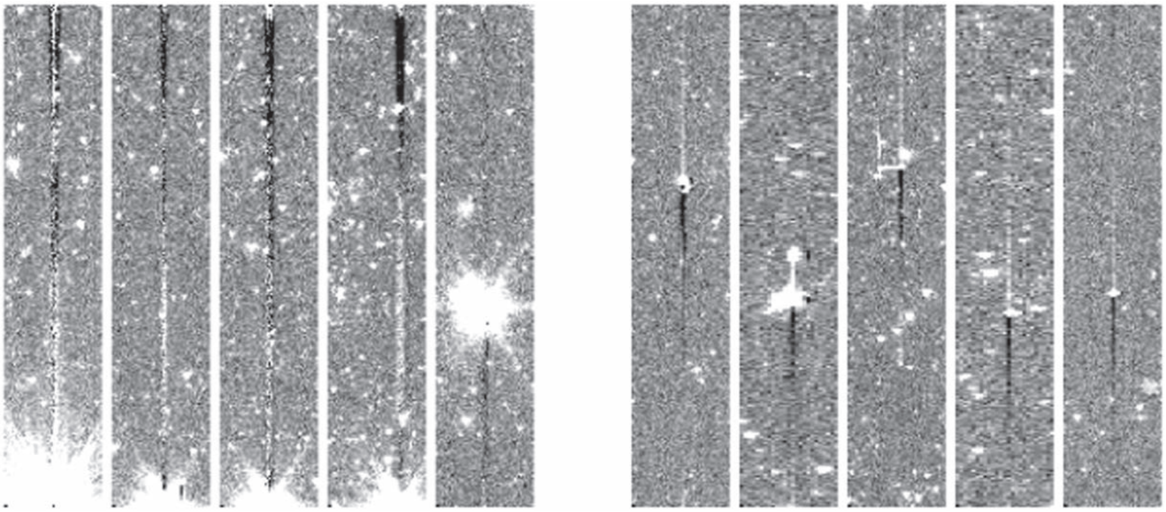

Figure 14. Examples of uncorrectable pull-down artifacts (left) and charged-particle strikes (right).

median value of the boundary region is always 2.0 MJy sr^{-1} or higher above the local background whereas for charged-particle strikes it is always well under 1.0 MJy sr^{-1}.

7.2. Pixel Mask Augmentation

The bad pixels in the channel 1 and channel 2 arrays are supplied as pixel masks in the set of standard CBCD pipeline data products for each AOR. For the time periods covered by the SMC-Last observations, these pixel masks are constant for both channels. However, it has been found through inspection of the mosaics and the contributing BCDs that, for channel 2, a number of pixels behave erratically at least over the course of an AOR, but are not flagged in the nominal pixel mask. While anomalous high values for a pixel are usually removed in the outlier rejection procedure, anomalous low values typically are not and result in "holes" in the mosaics.

These bad pixels are identified and added to the nominal pixel mask. For each AOR, the median value of each array pixel is calculated, giving the median array image for that AOR. Pixels that differ from the median of that image by at least 0.15 MJy sr^{-1} are taken to be bad pixels (the rms of the AOR-median images is typically less than 0.01 MJy sr^{-1}). This gives an ad hoc bad pixel mask for that AOR. In the generation of the mosaics, for each mosaic, a new bad pixel mask is created that is the union of the pixel masks for each AOR included in the mosaic and the nominal pipeline pixel mask. The nominal channel 2 pixel mask has 97 bad pixels flagged; the augmented pixel masks typically have 10–50 additional pixels flagged.

7.3. Anomalous Outlier Rejection

In the mosaic construction, we have used the outlier, dual outlier, and box outlier options available in the Mopex utility. We have found that these outlier rejection mechanisms together generally work well for the typical coverages in our data. However, this outlier rejection has an anomalous behavior in which, if a mosaic pixel is covered by data for two BCDs, and there are additionally two or more BCDs covering the mosaic pixel that have been masked at that location, the rejection mechanism tends to regard all local maxima above a very small amplitude, and the immediately surrounding regions, as outliers and deletes them from the mosaics. The cause of this anomaly is unclear, but we have applied a mitigation procedure separately for the single-epoch and combined-epoch mosaics.

For the single-epoch mosaics, sets of mosaics were created for both channels and both epochs with the outlier rejection mechanism turned off entirely. The standard mosaics were then compared to these specially prepared mosaics. Where the standard mosaic showed a data hole, as indicated by a not-a-number (NaN) value in the coverage image, and the prepared mosaic showed data present, this was considered a candidate for transferring the data from the prepared mosaic to the standard mosaic. An interactive routine, with a visual inspection of the regions from the two mosaics, was used to distinguish true cases of anomalous rejection from actual outliers. For the true cases, the data (for image, coverage, and uncertainty mosaics) were transferred, for a region including the NaN "hole" and also a two-pixel boundary surrounding the hole, because an anomalously rejected BCD pixel covering the hole can cause a data deficit for mosaic pixels up to 2 pixels away from the hole, given that the mosaic pixel size is half the angular size of the BCD pixels.

For the combined-epoch mosaics, the mosaics were similarly compared with these rejection-reprocessed single-epoch mosaics. Where a data hole in the combined-epoch mosaic corresponded to data present in the single-epoch mosaics, the data were transferred, also with a 2 pixel boundary around the hole. If data were present in one of the single-epoch mosaics but not both, the appropriate data were simply transferred to the combined-epoch mosaic. If data were present in both epochs, the image pixel values were calculated as the weighted average of the corresponding single-epoch pixel values, with the weights being the coverage values in the single-epoch mosaics (noting as in Section 3 that we are effectively using integration time as the weighting in the mosaics rather than the uncertainty values). The coverage is simply the sum of the individual coverages. The resulting uncertainty for the combined-epoch mosaic pixels is also a weighted combination. For uncertainties σ_1 and σ_2 for the two epochs, and coverages c_1 and c_2, the combined uncertainty is

$$\sigma_{\text{combined}} = \frac{\sqrt{c_1^2\sigma_1^2 + c_2^2\sigma_2^2}}{c_1 + c_2}. \quad (2)$$

7.4. Background Striations

Several of the mosaics show a distinctive striation pattern in the backgrounds. Figure 15 shows an example for epoch 1, channel 2. Where they occur, they occur to some extent in both channels and occur in both epochs, albeit in different locations. The amplitude is around 0.01–0.02 MJy sr^{-1}. The striations are not sufficiently regular to allow a correction, so they have been left unaddressed.

It appears that the striations are confined to a limited set of AORs, seven in epoch 1 and four in epoch 2, and originate from irregularities in the backgrounds in the BCDs, roughly horizontal patterns that persist for a few successive frames and then evolve to a different pattern. As the horizontal patterns are also parallel to the scan direction, this gives rise to the extended features in the mosaics.

Table 2 shows the affected AORs with the date and start UT.

8. Discussion and Conclusions

The SMC-Last program surveyed the SMC at 3.6 and 4.5 μm in two epochs, the first from 2017 August to September, and the second from 2017 November to 2018 February. We have created sets of 52 $1\overset{\circ}{.}1 \times 1\overset{\circ}{.}1$ images of the SMC from IRAC 3.6 and 4.5 μm observations. The data were corrected for instrumental artifacts and background discontinuities prior to mosaicking. The processed data result in six mosaics, three in each filter, with one from each of the two epochs and a third from the combined epochs. The final mosaics are available as FITS files from the Infrared Science Archive (IRSA).

The combined data sets from S^{3}MC, SAGE-SMC, SAGE-VAR, and SMC-LAST provide four epochs with the 3.6 and 4.5 μm IRAC filters covering the entire SMC and a fifth in the Bar, and four more in the center, which totals between four and nine epochs covering a temporal baseline of 9 yr everywhere in the SMC and 12 yr in the center. This temporal coverage opens a rich variety of research avenues ranging from astrometric studies in the foreground to variability studies in the SMC and the background.

Data from the Wide-field Infrared Survey Explorer (WISE; Wright et al. 2010), currently operating as the NEOWISE-R mission (Mainzer et al. 2014), provide even more temporal coverage. WISE observes the SMC every six months in two

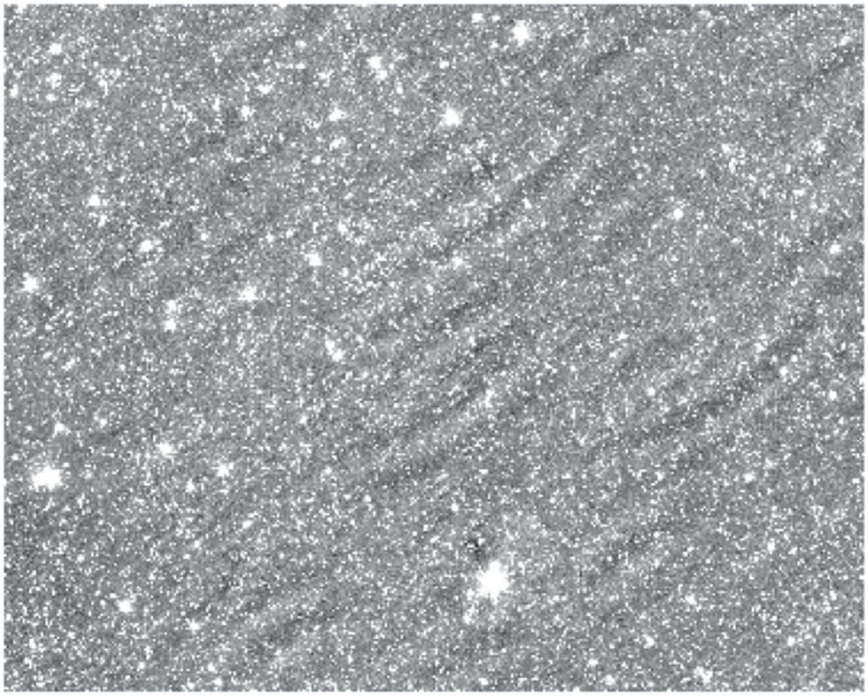

Figure 15. Example of striation pattern artifact in the mosaics. This is from the epoch 1, channel 2 data. The field is 1° wide and centered at 00^h52^m, $-70°54'$. The amplitude of the striations is about 0.01–0.02 MJy sr^{-1}. The "culprit" is AORKEY 64022016.

Table 2
AORs with Background Striations

AORKEY	Date			Start UT (hh:mm)
64018176	2017	Aug	25	00:19
64017920	2017	Aug	25	03:20
64020480	2017	Aug	29	02:10
64020224	2017	Aug	29	05:01
64019968	2017	Aug	29	07:52
64021760	2017	Sep	02	11:06
64022016	2017	Sep	05	08:52
65213952	2018	Jan	07	15:37
65215488	2018	Jan	14	12:42
65215744	2018	Jan	21	13:11
65255936	2018	Feb	08	17:03

filters with similar wavelengths to the 3.6 and 4.5 μm IRAC bands. These data can be combined with the IRAC data, although some differences must be considered. The WISE filters differ slightly from the IRAC filters and are centered at 3.4 and 4.6 μm. Their angular resolution is $\sim 6''$, compared to $\sim 2''$ for the IRAC data during the post-cryogenic mission (IRAC Instrument & Instrument Support Teams 2021, v. 4, their Table 2.1). The limiting magnitudes at 3.4 and 4.6 μm for the cryogenic WISE mission were 16.5 and 15.5 mag (respectively, Wright et al. 2010), and ~15.5 and 14.3 mag for the post-cryogenic mission (Mainzer et al. 2014) for single-epoch photometry. The SAGE-SMC project, performed during Spitzer's cryogenic phase, had corresponding limiting magnitudes of 17.6 and 17.0 mag (Gordon et al. 2011). The SMC-Last point-source extraction is still in progress (Kuchar et al. 2022, in preparation), but the preliminary estimates for the limiting magnitudes are ~17.1 and 16.8 mag. Figure 16 compares a small region of the SMC in the two 4 μm bands.

The WISE observations come at a steady cadence of one epoch every 180 days (with a gap from 2010 to the beginning of the NEOWISE-R mission in 2014). With the release of data from 2021, the extended WISE mission provides a total of 20 epochs. For targets that can be observed with WISE, the temporal coverage in the SMC now extends to between 13 and 16 yr, with additional epochs expected. One issue to be considered is the steady six-month cadence, which leads to potential aliasing problems. For example, studies of long-period variables can be hampered because 1 yr is close to the pulsation period of many Mira variables (Whitelock et al. 1994). The IRAC observations can break that degeneracy.

The SMC-Last data will build on the legacy of the previous studies of the infrared variability of long-period variables (LPVs) in the SMC (e.g., Riebel et al. 2015). Multiepoch optical studies have identified ~20,000 LPVs across the entire

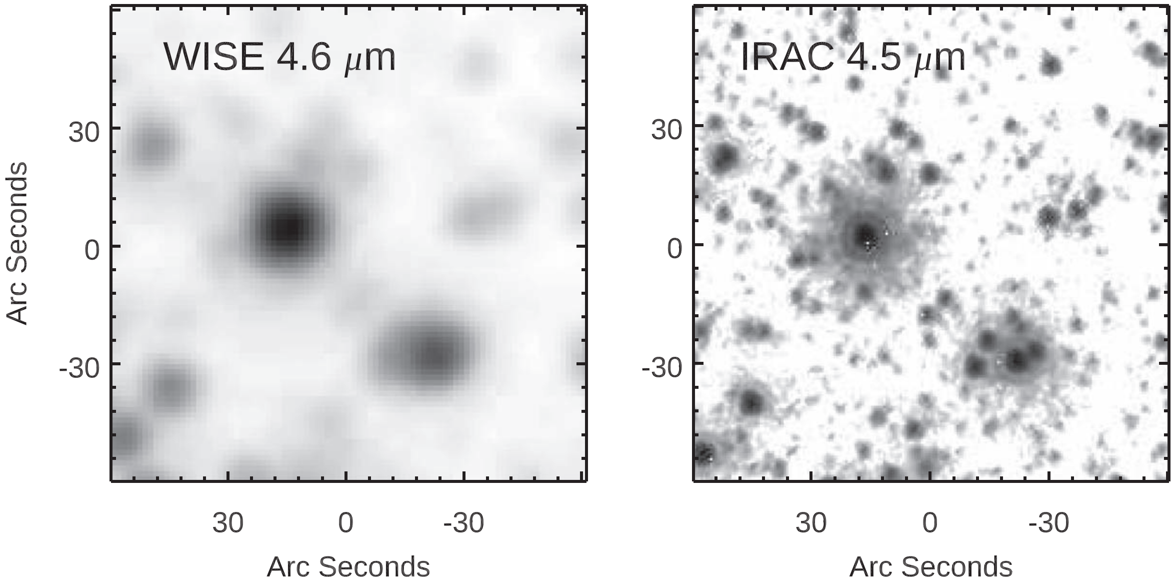

Figure 16. Comparison of WISE and IRAC, a 2′ box centered on $(\alpha, \delta) = (12.4121, -73.0397)$. (Left) WISE 4.6 μm; (right) IRAC 4.5 μm.

galaxy (e.g., Soszyński et al. 2011), but as stars die, they embed themselves in optically thick dust shells and become too faint for detection in the optical. These heavily enshrouded and optically invisible stars dominate the observed dust production in the SMC (Boyer et al. 2012; Srinivasan et al. 2016), but the optical multiepoch studies miss them. Similar objects in the Large Magellanic Cloud are barely at the edge of detection with WISE (Sloan et al. 2016; Groenewegen & Sloan 2018). To study them in the SMC, the greater sensitivity of SMC-Last and other Spitzer-based surveys is required.

The additional epochs from SMC-Last will also facilitate the study of young stellar objects (YSOs) in the SMC. The variability of T Tauri stars has been recognized for decades (Joy 1945), and more recent studies reveal that nearly all YSOs vary at some level (e.g., Megeath et al. 2012; Rice et al. 2012, 2015). The lower metallicity of the SMC leads to more luminous YSOs compared to the Galactic counterparts, and the lower dust abundances make them visible earlier in the evolution (de Wit et al. 2003). Most of the star-forming regions in the SMC were included in the S^{3}MC survey, so SMC-Last provides a 12 yr baseline to study this variability and search for transient events like FU Ori-like eruptions. SMC-Last will also enable searches for transients in the background population of galaxies.

Astrometry of the foreground population can reveal brown dwarfs due to their large proper motions. Assuming a limit of 17.0 mag at 4.5 μm, SMC-Last can detect early L dwarfs to 275 pc, early T dwarfs to 115 pc, and early Y dwarfs out to 25 pc (Kirkpatrick et al. 2011, 2012). We estimate, based on space densities from Cruz et al. (2003) and Kirkpatrick et al. (2012), that a comparison of SMC-Last to SAGE-SMC should reveal the motion of many dozens of L dwarfs, ∼10 T dwarfs, and 0.3 Y dwarfs in the field of 30 deg^2.

Thus, the two new epochs of infrared photometry from SMC-Last create many new temporal research opportunities in the SMC, behind the galaxy, and in the foreground.

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 NASA contract 1407. We particularly thank Nancy Silbermann of the Spitzer Science Center for her assistance in scheduling the "fill-in" observations. We have made use of NASA's Astrophysics Data System and CDS's Vizier service. Financial support for this work was provided by NASA through NASA ADAP grant 80NSSC19K0585. Finally, we thank the reviewer for the many helpful suggestions for changes and additions.

Appendix
Pointing Refinements

The mosaics were initially constructed using the "superboresight" pointing solution, i.e., the nominal astrometry information in the BCD headers, which the IRAC instrument handbook indicates should have an rms accuracy of about 0.″16. However, in analyzing source extractions from the initial mosaics, it was found that the measured source-position differences between the two epochs, particularly the R.A., typically vary in curiously systematic patterns. Figure 17 shows the R.A. differences between epoch 1 and epoch 2 channel 1 sources for the mosaic centered at $01^{h}31^{m}$, −73°22′. For this image, the R.A. differences are binned in 2.′0 bins, and the bin

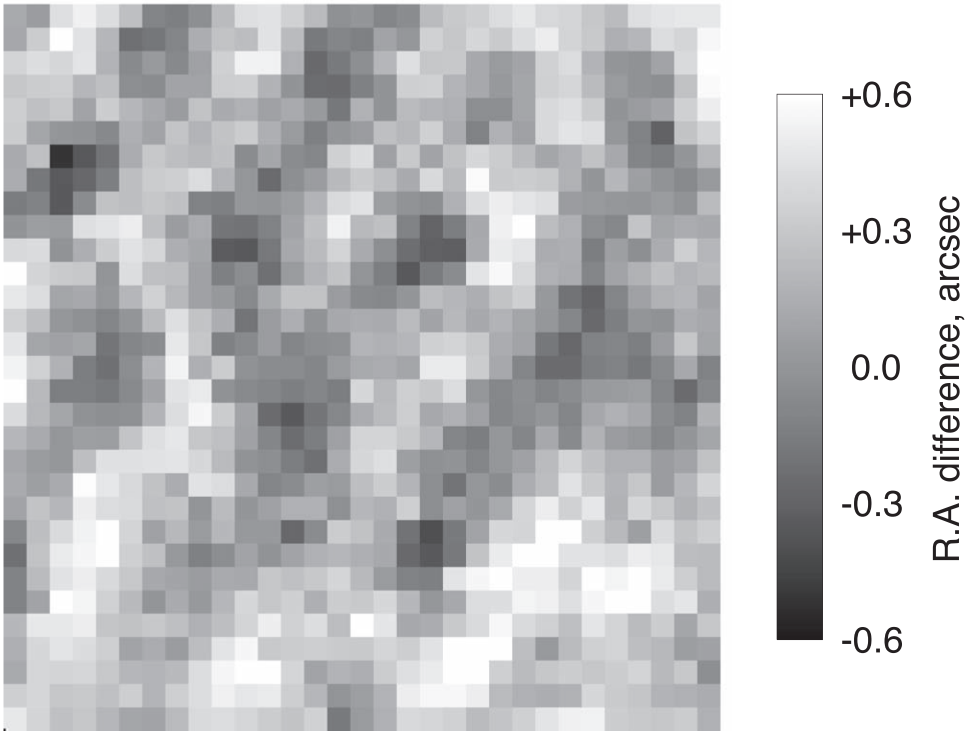

Figure 17. R.A. differences between epoch 1 and epoch 2 sources, with the "superboresight" pointing, for the ∼1°.1 × 1°.1 mosaic centered at $01^{\mathrm{h}}31^{\mathrm{m}}$, −73°22′. The sources have been binned in 2′ bins and the image values calculated as the median of the R.A. differences in the bin. This is channel 1 data, but channel 2 shows a nearly identical pattern.

values are the median difference values in the bin. Most mosaic regions show a generally similar pattern in the R.A., to varying degrees in both epochs, but not for decl. with the exception noted below. The patterns for channel 1 and 2 are nearly identical.

The cause of these patterns has been found to be quasi-periodic errors in the center R.A. (i.e., CRVAL1 keyword) in the BCD headers. Figure 18 shows the errors for the first 150 BCDs for AORKEY 64019968, channel 1. The absolute errors have been calculated by mapping 2MASS 6x sources (Cutri et al. 2012) into the astrometry for each BCD, matching to the sources in the BCD image, and taking the median of the differences in R.A. and decl. over all the matched sources. As noted above, the R.A. errors for channel 2 are nearly identical.

All of our AORs but six include the "superboresight" pointing. The six exceptions are epoch 2 AORs covering the SMC core. For these AORs, both the R.A. and decl. errors show similar periodicity and amplitude, and generally have larger global pointing offset errors as well, up to about 0″.5.

Most of our CBCD data, as provided by the Spitzer Science Center, include the "refined" pointing solution, for which the astrometry of each BCD has been corrected to absolute coordinates by matching the sources on the BCD frame to sources from the 2MASS catalog, and resetting the parameters (CRVAL1, CRVAL2, and the CD matrix) to minimize the differences. For our data, the "refined" pointing solution shows a significant improvement over the superboresight pointing. However, a small quantity of our data is lacking the refined solution, so we have developed a procedure to correct the absolute coordinates for each BCD using the 2MASS 6x catalog and applied it to all our data. The details of the

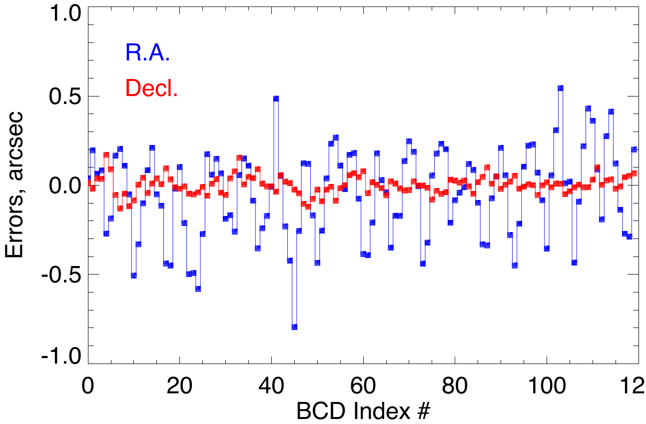

Figure 18. R.A. and decl. errors for the first 120 BCDs of AORKEY = 64019968, channel 1, using the "superboresight" pointing solution, showing the quasi-periodic error in R.A. The errors have been calculated by matching 2MASS 6x sources to sources on each BCD and taking the median difference over the matched sources. The spacing between BCDs is about 26 s.

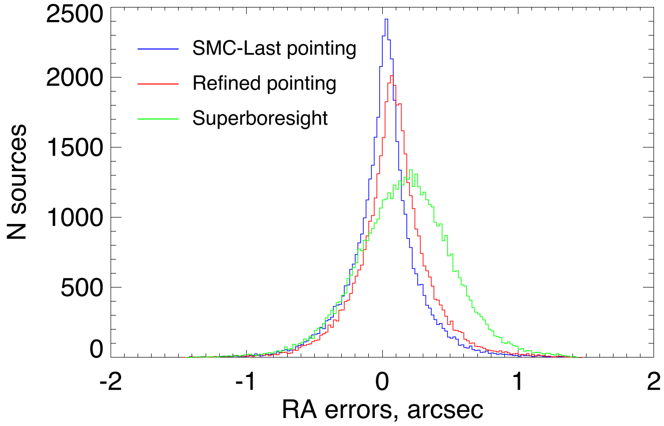

Figure 19. Histogram of R.A. errors for the SMC-Last pointing solution and the “refined” solution, for the mosaic centered at $01^{\rm h}07^{\rm m}$, $-72°22'$, for epoch 2, channel 1. The errors are calculated by matching the sources extracted from the mosaic with the 2MASS 6x sources. The “superboresight” data included for comparison are from epoch 1 data for this mosaic because the epoch 2 data for this mosaic did not include that pointing solution.

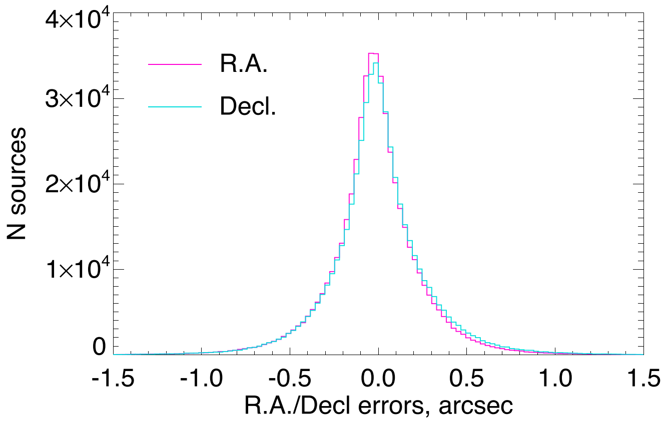

Figure 20. Histogram of R.A. and decl. errors for the SMC-Last pointing solution for all epoch 1, channel 1 sources matched to 2MASS 6x (about 510,000 sources). The other epoch and channel results are similar. The median errors are all under $0\overset{''}{.}03$, and the rms values are $\sim 0\overset{''}{.}25$.

procedure are presented in Kuchar T. A. et al. (2022, in preparation). Figure 19 shows the results of source extractions from mosaics created using our pointing correction procedure, and from the “refined” solution, with errors determined from matching to the 2MASS 6x catalog sources. Also shown in the figure for comparison are results for the “superboresight” pointing.

Figure 20 shows both R.A. and decl. errors for the results of our pointing solution for all of the epoch 1, channel 1 mosaics.

ORCID iDs

D. R. Mizuno, https://orcid.org/0000-0003-0947-2824
Kathleen E. Kraemer https://orcid.org/0000-0002-2626-7155
T. A. Kuchar https://orcid.org/0000-0003-1955-8509
G. C. Sloan, https://orcid.org/0000-0003-4520-1044

References

Bolatto, A. D., Simon, J. D., Stanimirović, S., et al. 2007, ApJ, 655, 212
Boyer, M. L., Srinivasan, S., Riebel, D., et al. 2012, ApJ, 748, 40
Choudhury, S., de Grijs, R., Rubele, S., et al. 2020, MNRAS, 497, 3746
Cioni, M.-R. L., Clementini, G., Girardi, L., et al. 2011, A&A, 527, A116
Cruz, K. L., Reid, I. N., Liebert, J., Kirkpatrick, J. D., & Lowrance, P. J. 2003, AJ, 126, 2421
Cutri, R. M., Skrutskie, M. F., van Dyk, S., et al. 2012, yCat, 2281 II/281
de Grijs, R., & Bono, G. 2015, AJ, 149, 179
de Wit, W. J., Beaulieu, J. P., Lamers, H. J. G. L. M., Lesquoy, E., & Marquette, J. B. 2003, A&A, 410, 199
Fazio, G. G., Hora, J. L., Allen, L. E., et al. 2004, ApJS, 154, 10
Gordon, K. D., Meixner, M., Meade, M. R., et al. 2011, AJ, 142, 102
Graczyk, D., Pietrzyński, G., Thompson, I. B., et al. 2020, ApJ, 904, 13
Groenewegen, M. A. T., & Sloan, G. C. 2018, A&A, 609, A114
Hora, J. L., Patten, B. M., Fazio, G. G., & Glaccum, W. J. 2006, Proc. SPIE, 6276, 62760J
Hora, J. L., Fazio, G. G., Allen, L. E., et al. 2004, Proc. SPIE, 5487, 77
Hora, J. L., Carey, S., Surace, J., et al. 2008, PASP, 120, 1233
IRAC Instrument and Instrument Support Teams 2021, IRAC Instrument Handbook, Tech. Rep., Infrared Science Archive, doi:10.26131/irsa486
Ita, Y., Matsunaga, N., Tanabé, T., et al. 2018, MNRAS, 481, 4206
Joy, A. H. 1945, ApJ, 102, 168
Kirkpatrick, J. D., Cushing, M. C., Gelino, C. R., et al. 2011, ApJS, 197, 19
Kirkpatrick, J. D., Gelino, C. R., Cushing, M. C., et al. 2012, ApJ, 753, 156
Luck, R. E., Moffett, T. J., Barnes, T. G. I., & Gieren, W. P. 1998, AJ, 115, 605
Mainzer, A., Bauer, J., Cutri, R. M., et al. 2014, ApJ, 792, 30
Makovoz, D., & Khan, I. 2005, in ASP Conf. Ser., 347, Astronomical Data Analysis Software and Systems XIV, ed. P. Shopbell, M. Britton, & R. Ebert (San Francisco, CA: ASP), 81
Megeath, S. T., Gutermuth, R., Muzerolle, J., et al. 2012, AJ, 144, 192
Meixner, M., Gordon, K. D., Indebetouw, R., et al. 2006, AJ, 132, 2268
Mizuno, D. R. 2008, AFRL_BCD_OVERLAP, Contributed software, Spitzer Science Center, http://irsa.ipac.caltech.edu/data/SPITZER/docs/dataanalysistools/tools/contributed/irac/afrlbcdoverlap/
Mizuno, D. R., Carey, S. J., Noriega-Crespo, A., et al. 2008, PASP, 120, 1028
Nidever, D. L., Monachesi, A., Bell, E. F., et al. 2013, ApJ, 779, 145
Rice, T. S., Reipurth, B., Wolk, S. J., Vaz, L. P., & Cross, N. J. G. 2015, AJ, 150, 132
Rice, T. S., Wolk, S. J., & Aspin, C. 2012, ApJ, 755, 65
Riebel, D., Boyer, M. L., Srinivasan, S., et al. 2015, ApJ, 807, 1
Rieke, G. H., Young, E. T., Engelbracht, C. W., et al. 2004, ApJS, 154, 25
Rubele, S., Pastorelli, G., Girardi, L., et al. 2018, MNRAS, 478, 5017
Scowcroft, V., Freedman, W. L., Madore, B. F., et al. 2016, ApJ, 816, 49
Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, AJ, 131, 1163
Sloan, G. C., Kraemer, K. E., McDonald, I., et al. 2016, ApJ, 826, 44
Soszyński, I., Udalski, A., Szymański, M. K., et al. 2011, AcA, 61, 217
Srinivasan, S., Boyer, M. L., Kemper, F., et al. 2016, MNRAS, 457, 2814
Subramanian, S., Rubele, S., Sun, N.-C., et al. 2017, MNRAS, 467, 2980
Udalski, A., Kubiak, M., & Szymanski, M. 1997, AcA, 47, 319
Udalski, A., Soszyński, I., Szymański, M. K., et al. 2008, AcA, 58, 329
Whitelock, P., Menzies, J., Feast, M., et al. 1994, MNRAS, 267, 711
Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, AJ, 140, 1868
Yanchulova Merica-Jones, P., Sandstrom, K. M., Johnson, L. C., et al. 2021, ApJ, 907, 50