

SMC-Last Mosaic Images

D. R. MIZUNO,¹ 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*

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*

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 to 2017 September 13, and 2017 November 24 to 2018 February 12. The survey region comprises ~ 30 square degrees covering the SMC and the Bridge to the Large Magellanic Cloud. The region is covered by 52 $\sim 1^\circ 1' \times 1^\circ 1'$ 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 $0''.6$ with a resolution of $\sim 2''.0$. We describe processing to correct or mitigate residual artifacts and remove background discontinuities. The mosaic images are publically available at the Infrared Science Archive (IRSA).

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 kpc 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 $[\text{Fe}/\text{H}] = -0.68$ (Luck et al. 1998), while for the red-giant population, the mean $[\text{Fe}/\text{H}] = -0.97$ (Choudhury et al. 2020). Rubele et al. (2018) find that $[\text{Fe}/\text{H}]$ 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 square degrees 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 square degrees 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 (S³MC). The survey covered roughly 10 square degrees 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. 2004a) 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. S³MC was 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 square degree 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. 2011a). 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 square degrees 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 square degree 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 nine years 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 years.

The combined data sets provide leverage on the variability of stars in the SMC, both young and old. The 12-year baseline will help identify FU Ori-type transients among the young stellar population within the SMC and even transients in the galaxies behind it. When combined with data from WISE (Wright et al. 2010) and the reactivated mission, NEOWISE-R (Mainzer et al. 2014), the Spitzer epochs can break the degeneracies associated with the steady 180-day cadence at which WISE observes the SMC, and could even make it possible to fit periods and amplitudes to long-period variables too deeply embedded in their own dust shells to be observed at shorter wavelengths (e.g., Sloan et al. 2016; Groenewegen & Sloan 2018). The 12-year baseline and the sensitivity of Spitzer also may make it possible to identify brown dwarfs in the foreground from their proper motions.

2. OBSERVATIONS

We mapped ~ 30 square degrees 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. 2004b). The Astronomical Observation Requests (AORs) were based on those used for the 2008 maps for SAGE-SMC from Gordon et al. (2011b). The observations

were taken in the high-dynamic range mode (HDR), with 0.4 and 10.4 second 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 $\sim 1^\circ 1' \times 1^\circ 1'$. 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 two in coverage in a single map, enabling 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 to 2017 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 lasted 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. IMAGE PRODUCTS OVERVIEW

The SMC-Last mosaics comprise, for each epoch and channel, 52 FITS image files, each covering a region approximately $1^\circ 06' \times 1^\circ 06'$, with an angular pixel size of $0''.6$ and a resolution of $\sim 2''.0$. 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 declination (Decl.), and in each row the images are approximately 1° apart in right ascension (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 $4.5 \mu\text{m}$ 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.

The mosaics have been constructed using the mosaic utility in the Mopex software package provided by the Spitzer Science Center ([Makovoz & Khan 2005](#)). We note two peculiarities with this utility: the first is that the data uncertainties are not used for weighting the averaging in the mosaic pixel filling, and so combining data with significantly different uncertainties is problematic (see [section 7.1](#)); and second, the outlier rejection mechanism has erroneous behavior under certain circumstances (see [section 7.3](#)).

The Corrected Basic Calibrated Data (CBCD) pipeline products provided by the Spitzer Science Center have been used as the input data source to construct the mosaics. The supplied CBCD

data products have been corrected for detector linearization and flatfielding, calibrated to flux units (MJy sr^{-1}), have had dark current and a zodiacal model removed, have had the pointing refinement applied, and have been flagged for radiation hits, saturation effects, latent images, stray light and other artifacts (IRAC Instrument and Instrument Support Teams 2021). The 10.4s integrations only have been used in the mosaics, except where, as noted below (section 7.1), the 0.4s integration data are used to fill in saturated pixels for very bright sources.

The CBCD products have been processed further prior to the construction of the mosaics, as described in the remainder of this paper. Section 4 describes adjustment of background levels, largely to remove discontinuities between AORs; Section 5 describes corrections of column pull-down/pull-up artifacts and latent images; Section 6 describes a more thorough masking of uncorrectable artifacts than the pipeline provides; and Section 7 describes fill-in of saturated sources and mitigation of the outlier rejection anomaly.

In the following, the term “pipeline” refers to the CBCD processing pipeline, and “BCD” refers to a single frame of 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 non-astronomical 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 in pixel response 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-hour 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 3 or 4 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. However, we have found that the RMS noise levels in the BCDs actually decrease slightly, by about 5-10%, indicating that the “noise” in the correction images is true pixel-to-pixel variations in response.

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:

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

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 the process is repeated 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 5 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 more quiescent region. 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 $\sim 1^\circ$ 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 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 are not going to change observably over the timescale of the two epochs.

Figure 6 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 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, although the alternative is to postulate that there is some systematic elevated bias in the raw data for the SMC core that roughly scales with the true background levels.

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.

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 very quiescent, 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 non-elevated background levels over the rest of the BCD frame, for the bandlike 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 7 shows a few examples the pull-down artifacts and the resulting corrections, and Figure 8 likewise for pull-up artifacts.

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 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 9 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.

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.

Figure 9 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 9 BCD window median image is taken for the whole array, and, again exploiting the generally quiescent 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 10 shows a few examples of stray light, both caught and missed by the CBCD pipeline.

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. Towards 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, 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 otherwise have no apparent pull-down effect, in spite of being flagged as such. 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 11 shows an example of a column with this artifact and the uncorrected effect on the resulting mosaic image.

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, are 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 12 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 substitution of the corresponding pixel values for the uncertainty data. (The pointing from the 0.4s to the 10.4s integrations seems to be sufficiently stable for the pixel-for-pixel substitution.)

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 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*

It has been found that the Mopex mosaic construction utility has an anomalous behaviour 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 which have been masked at that location, the outlier 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 the 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 a BCD pixel anomalously rejected covering the hole can cause a data deficit for mosaic pixels up to two 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 two-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 (note that the Mopex mosaic facility does not use the uncertainty values in averaging the BCD pixels to determine the mosaic pixel 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:

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

7.4. *Background striations*

Several of the mosaics show a distinctive striation pattern in the backgrounds. Figure 13 shows an example for epoch 1, channel 2. Where they occur, they occur to some extent in both channels, and




Figure 1. Final coverage of the SMC-Last program. (left) epoch 1, 2017 Aug 25 - Sep 13; (right) epoch 2, 2017 Nov 24 - 2018 Feb 12, overlaid on an IRAS 12 μm image. The SAGE-SMC coverage is shown as the cyan outline. Blue circles in the righthand panel show locations of IRS observations.

occur in both epochs, albeit in different locations. The amplitude is around $0.01\text{-}0.02 \text{ MJy sr}^{-1}$. 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.

8. SUMMARY

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 $1^\circ 1' \times 1^\circ 1'$ 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).

ACKNOWLEDGMENTS

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. Financial support for this work was provided by NASA through NASA ADAP grant 80NSSC19K0585.




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

REFERENCES

- Bolatto, A. D., Simon, J. D., Stanimirović, S., et al. 2007, *ApJ*, 655, 212, doi: [10.1086/509104](https://doi.org/10.1086/509104)
- Choudhury, S., de Grijs, R., Rubele, S., et al. 2020, *MNRAS*, 497, 3746, doi: [10.1093/mnras/staa2140](https://doi.org/10.1093/mnras/staa2140)
- Cioni, M.-R. L., Clementini, G., Girardi, L., et al. 2011, *A&A*, 527, A116, doi: [10.1051/0004-6361/201016137](https://doi.org/10.1051/0004-6361/201016137)
- de Grijs, R., & Bono, G. 2015, *AJ*, 149, 179, doi: [10.1088/0004-6256/149/6/179](https://doi.org/10.1088/0004-6256/149/6/179)
- Fazio, G. G., Hora, J. L., Allen, L. E., et al. 2004a, *ApJS*, 154, 10, doi: [10.1086/422843](https://doi.org/10.1086/422843)
- . 2004b, *ApJS*, 154, 10, doi: [10.1086/422843](https://doi.org/10.1086/422843)
- Gordon, K. D., Meixner, M., Meade, M. R., et al. 2011a, *AJ*, 142, 102, doi: [10.1088/0004-6256/142/4/102](https://doi.org/10.1088/0004-6256/142/4/102)
- . 2011b, *AJ*, 142, 102, doi: [10.1088/0004-6256/142/4/102](https://doi.org/10.1088/0004-6256/142/4/102)
- Graczyk, D., Pietrzyński, G., Thompson, I. B., et al. 2020, *ApJ*, 904, 13, doi: [10.3847/1538-4357/abbb2b](https://doi.org/10.3847/1538-4357/abbb2b)
- Groenewegen, M. A. T., & Sloan, G. C. 2018, *A&A*, 609, A114, doi: [10.1051/0004-6361/201731089](https://doi.org/10.1051/0004-6361/201731089)
- Hora, J. L., Patten, B. M., Fazio, G. G., & Glaccum, W. J. 2006, in *Proc. SPIE*, Vol. 6276, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 62760J, doi: [10.1117/12.672017](https://doi.org/10.1117/12.672017)
- Hora, J. L., Fazio, G. G., Allen, L. E., et al. 2004, in *Proc. SPIE*, Vol. 5487, Optical, Infrared, and Millimeter Space Telescopes, ed. J. C. Mather, 77–92, doi: [10.1117/12.550744](https://doi.org/10.1117/12.550744)
- Hora, J. L., Carey, S., Surace, J., et al. 2008, *PASP*, 120, 1233, doi: [10.1086/593217](https://doi.org/10.1086/593217)
- IRAC Instrument and Instrument Support Teams. 2021, *IRAC Instrument Handbook*, Tech. rep., Infrared Science Archive, doi: [10.26131/irsa486](https://doi.org/10.26131/irsa486)




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.




Figure 5. 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 more quiescent region.




Figure 6. 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.




Figure 7. 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 8. Examples of the pull-up artifact and their corrections. The bright source in the band-like artifact at 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.




Figure 9. 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 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.




Figure 10. Some examples of the stray light artifact. The top row are cases caught by the CBCD pipeline, and the remainder were missed. Note that the pipeline misses some very bright cases.

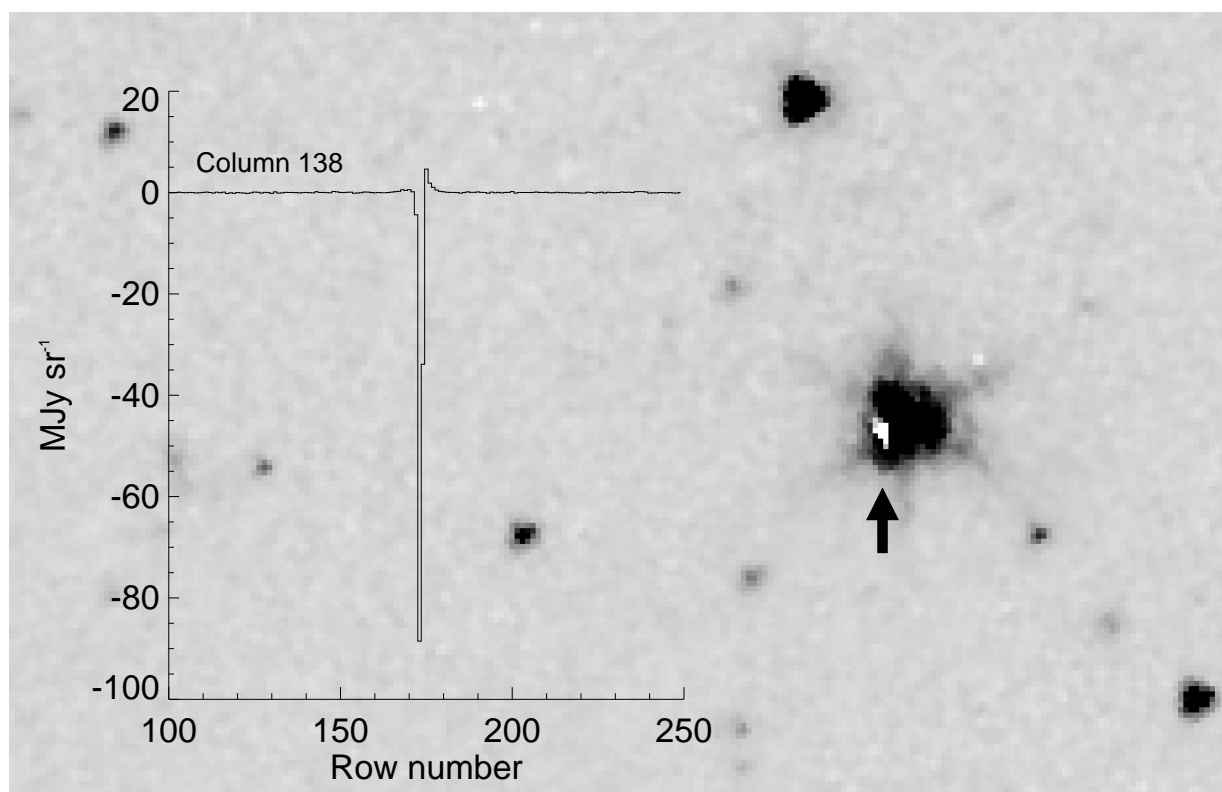


Figure 11. Example of the very-negative pixel values seen in some “pull-down” columns, and their effect on resulting mosaic images. Note that this is not a data hole in the mosaic, rather a very negative value, because the outlier rejection mechanism fails to trap these cases. The affected source is at $\sim 02^h 00^m, -74^\circ 52'$.




Figure 12. Examples of uncorrectable pull-down artifacts (left) and charged-particle strikes (right).




Figure 13. 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 $00^h52^m, -70^\circ54'$. The amplitude of the striations is about $0.01\text{--}0.02$ MJy sr^{-1} . The “culprit” is AORKEY 64022016.

Ita, Y., Matsunaga, N., Tanabé, T., et al. 2018, MNRAS, 481, 4206, doi: [10.1093/mnras/sty2539](https://doi.org/10.1093/mnras/sty2539)

Luck, R. E., Moffett, T. J., Barnes, Thomas G., I., & Gieren, W. P. 1998, AJ, 115, 605, doi: [10.1086/300227](https://doi.org/10.1086/300227)

Mainzer, A., Bauer, J., Cutri, R. M., et al. 2014, ApJ, 792, 30, doi: [10.1088/0004-637X/792/1/30](https://doi.org/10.1088/0004-637X/792/1/30)

Makovoz, D., & Khan, I. 2005, in Astronomical Society of the Pacific Conference Series, Vol. 347, Astronomical Data Analysis Software and Systems XIV, ed. P. Shopbell, M. Britton, & R. Ebert, 81

Meixner, M., Gordon, K. D., Indebetouw, R., et al. 2006, AJ, 132, 2268, doi: [10.1086/508185](https://doi.org/10.1086/508185)

Mizuno, D. R. 2008, AFRL_BCD_OVERLAP, Contributed software, Spitzer Science Center, irsa.ipac.caltech.edu/data/SPITZER/docs/dataanalysistools/tools/contributed/irac/afrlbcdoverlap/.

<http://irsa.ipac.caltech.edu/data/SPITZER/docs/dataana>

Mizuno, D. R., Carey, S. J., Noriega-Crespo, A., et al. 2008, PASP, 120, 1028, doi: [10.1086/591809](https://doi.org/10.1086/591809)

Nidever, D. L., Monachesi, A., Bell, E. F., et al. 2013, ApJ, 779, 145, doi: [10.1088/0004-637X/779/2/145](https://doi.org/10.1088/0004-637X/779/2/145)

Riebel, D., Boyer, M. L., Srinivasan, S., et al. 2015, ApJ, 807, 1, doi: [10.1088/0004-637X/807/1/1](https://doi.org/10.1088/0004-637X/807/1/1)

Rieke, G. H., Young, E. T., Engelbracht, C. W., et al. 2004, ApJS, 154, 25, doi: [10.1086/422717](https://doi.org/10.1086/422717)

- Rubele, S., Pastorelli, G., Girardi, L., et al. 2018, MNRAS, 478, 5017, doi: [10.1093/mnras/sty1279](https://doi.org/10.1093/mnras/sty1279)
- Scowcroft, V., Freedman, W. L., Madore, B. F., et al. 2016, ApJ, 816, 49, doi: [10.3847/0004-637X/816/2/49](https://doi.org/10.3847/0004-637X/816/2/49)
- Sloan, G. C., Kraemer, K. E., McDonald, I., et al. 2016, ApJ, 826, 44, doi: [10.3847/0004-637X/826/1/44](https://doi.org/10.3847/0004-637X/826/1/44)
- Soszyński, I., Udalski, A., Szymański, M. K., et al. 2011, AcA, 61, 217. <https://arxiv.org/abs/1109.1143>
- Subramanian, S., Rubele, S., Sun, N.-C., et al. 2017, MNRAS, 467, 2980, doi: [10.1093/mnras/stx205](https://doi.org/10.1093/mnras/stx205)
- Udalski, A., Kubiak, M., & Szymanski, M. 1997, AcA, 47, 319. <https://arxiv.org/abs/astro-ph/9710091>
- Udalski, A., Soszyński, I., Szymański, M. K., et al. 2008, AcA, 58, 329. <https://arxiv.org/abs/0901.4632>
- Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, AJ, 140, 1868, doi: [10.1088/0004-6256/140/6/1868](https://doi.org/10.1088/0004-6256/140/6/1868)
- Yanchulova Merica-Jones, P., Sandstrom, K. M., Johnson, L. C., et al. 2021, ApJ, 907, 50, doi: [10.3847/1538-4357/abc48b](https://doi.org/10.3847/1538-4357/abc48b)