

A Robust Algorithm for the Pointing Refinement and Registration of Astronomical Images

FRANK J. MASCI, DAVID MAKOVZ, AND MEHRDAD MOSHIR

Spitzer Science Center, California Institute of Technology, Pasadena, CA 91125; fmasci@ipac.caltech.edu

Received 2004 June 22; accepted 2004 July 15; published 2004 August 26

ABSTRACT. We present a generic algorithm for performing astronomical image registration and pointing refinement. The method is based on the matching of positions and fluxes of available point sources in image overlap regions. This information is used to compute a set of image offset corrections by globally minimizing a weighted sum of all matched point-source positional differences in a prespecified reference image frame. A fast linear sparse matrix solver is used for the minimization. From these corrections, the pointings and orientations of images can be refined in either a relative sense in which pointings become fixed (registered) relative to a single input image, or in an absolute sense (in the International Celestial Reference System [ICRS]) if absolute point-source information is known. The latter provides absolute pointing refinement to an accuracy that depends on the robustness of point-source extractions, match statistics, and accuracy of the astrometric catalog used. The software is currently used in the *Spitzer* image-processing pipelines, although it is adaptable to any astronomical imaging system that uses the FITS image format and world coordinate system (WCS) pointing standard. We test the algorithm using Monte Carlo simulations and compare them to image data acquired with the Infrared Array Camera (IRAC) onboard the *Spitzer Space Telescope*. We find that dispersions in matched source separations after refinement are entirely consistent with centroiding errors in source extractions, implying that systematic uncertainties due to inaccurately calibrated distortions are negligible. For these data, we predict refinements to better than ~ 70 and ~ 280 mas (2σ radial) for the IRAC 3.6 and $8\ \mu\text{m}$ bands, respectively. These bands bracket two extremes in available source matches, and for the data under study, correspond to an average of about 55 and 8 matches per image in these two bands, respectively.

1. INTRODUCTION

The ability to register or co-align images in astronomy to an accuracy better than the Nyquist sampling density of a detector's point-response function is pivotal if one wants to maximize the resolution and signal-to-noise ratio (S/N) in a mosaic of co-added images. Absolute accuracy is inhibited by instabilities in telescope pointing control, tracking sensors, and how these behave in the science instrument frame in the presence of thermomechanical disturbances. For instance, imaging detectors on both the *Hubble*¹ and *Spitzer*² *Space Telescopes* provide an absolute pointing accuracy of $\sim 0''.5$ – $1''$ (1σ radial). By comparison, the highest attainable resolutions³ are $\sim 0''.02$ and $1''$ for these telescopes, respectively. Without further refinement, the current pointing accuracies are insufficient to exploit the near-diffraction-limited resolution capabilities the detectors can provide. A factor of 10–20 improvement in pointing is required for optimal image registration.

Good image registration enables source extraction and position determination to fainter flux levels for a given S/N. Com-

parison or registration with astrometric sources whose positions are known to better than a few percent of the observed image pointing uncertainty also allows the refinement of image-frame pointings in the International Celestial Reference System (ICRS). Absolute pointing refinement can be achieved to an accuracy approaching that of the astrometric catalog used, or better, given good match statistics. This alleviates possible ambiguities when performing cross-identification/correlation of extracted sources across wavelength-dependent catalogs. Furthermore, the accurate placement of slits for follow-up spectroscopic studies requires source positions that are accurate to better than a few tenths of an arcsecond in the ICRS, or until the desired positional accuracy is limited by centroiding error in the array frame.

Broadly speaking, image registration methods can be loosely divided into three classes: algorithms that use information in pixel space directly (e.g., by correlating common objects; Barnea & Silverman 1972), algorithms that attempt to match features or identified parts of objects (known as “graph-theoretic” methods; Brown 1992), and algorithms that use the frequency domain (e.g., methods based on computing cross-correlation power spectra via the fast Fourier transform; Kuglin & Hines 1975). The two conventional methods for registering images in astronomy involve either interactively identifying common point

¹ See <http://www.stsci.edu/hst/observatory/pointing>.

² See <http://ssc.spitzer.caltech.edu/documents/SOM>.

³ Measured in terms of the full width at half-maximum (FWHM) of the center of an Airy disk pattern.

sources in overlapping image fields (i.e., object correlation), or using actual detector-acquired pointings (with inherent uncertainties) directly and estimating relative image offsets from them. These methods are available in most data-reduction packages (e.g., IRAF, Davis 1996; STARLINK, Bly et al. 2002) and are mostly limited by accuracies in point-source centroids, sufficient match statistics, telescope pointing control, or are subject to random-walk (cumulative) uncertainties.

A number of robust methods for astrometric calibration of single images have also been implemented in the commonly used data-reduction packages (Veran & Rigaut 1998; Valdes 1998; Bustos & Calderón 2000), although an automated, self-consistent means for simultaneous registration and refinement of multiple astronomical images comprising a mosaicked region is generally lacking. The single-image methods assume one has sufficient numbers of astrometric matches to mitigate against uncertainties in the centroids of extracted sources. We have gone a step further by using the available point-source content to obtain a global refined solution for all frames such that all matched point-source positional discrepancies in all frame-to-frame and frame-to-absolute overlaps are minimized. Combining the relative and absolute source information reduces the demand of having sufficient astrometric matches. This becomes important for mid- to far-infrared imagery, in which cross-wavelength astrometric calibration is often unreliable, because of differing sensitivities, source populations, and detector point-spread function (PSF) sizes.

We have developed an algorithm to simultaneously register and refine the pointing of an ensemble of astronomical images to accuracies better than that inherent in point-source centroid uncertainties (and dictated by point source match statistics). This paper describes the global minimization algorithm and presents a case study using data from *Spitzer's* Infrared Array Camera (IRAC). The outline is as follows. Section 2 describes the algorithm and the expected pointing refinement accuracies. Section 3 validates the algorithm, using a Monte Carlo simulation of IRAC data. Section 4 compares these results to real observations acquired with IRAC. Discussion and conclusions are given in § 5.

2. ALGORITHM

The algorithm has been implemented as a stand-alone software package called POINTINGREFINE.⁴ The main inputs to POINTINGREFINE are a FITS image list, with each FITS image containing the standard world coordinate system (WCS) keywords (Greisen & Calabretta 2002; Calabretta & Greisen 2002), an accompanying list of flux-calibrated point-source extraction tables, an optional astrometric source list, and configuration parameters. The software accepts point-source extraction tables that adhere to the format generated by the *Spitzer*

Science Center (SSC) source extractor. This includes programs that perform point-source extraction as well, although these will not be described here. A general overview of the processing steps involved in the global minimization algorithm (the heart of POINTINGREFINE) is shown in Figure 1. In the following sections, we expand on some of the more important computational steps of this algorithm.

2.1. Setup and Point-Source Matching

Prior to source matching, we first compute the total number of image pairs expected to be overlapping (which could potentially contain correlated/common sources) in the input ensemble of images. This is used for efficient a priori memory allocation. Given a number of images N_{imgs} , the maximum number of distinct frame pairs that can mutually overlap is

$$N_{\text{maxpairs}} = \frac{1}{2}N_{\text{imgs}}(N_{\text{imgs}} - 1). \quad (1)$$

This maximum occurs when all images are stacked more or less on top of each other. For a sparse mosaic, this number is smaller and thus puts less of a burden on the required memory. The total number of potential overlapping frame pairs is found by computing the distances between the centers of every image pair and finding whether the distance is less than the typical radius of a circle inscribing an image.

Source positions from all extraction tables (including absolute astrometric references, if specified) are sorted by declination. This precondition speeds up the source-matching procedure, converting it from an $O(N^2)$ to an approximate $O(N)$ linear process. Every possible overlapping image pair is scanned for common point sources in the R.A., decl. coordinate system. Both position and flux matching is performed. The position-match step attempts to find sources that fall within a nominal search radius (typically several times the root sum squares of prior image pointing uncertainty and extraction centroid error). If more than one match is found within the search radius, no match is declared, because of the possible ambiguity. In other words, only singly matched sources within the search radius are used. Sources are also simultaneously matched in flux. A flux match is satisfied if two fluxes fall within a maximum tolerable relative flux difference threshold. The software includes options for rescaling the input fluxes of astrometric references to be commensurate with those of actual extractions. A minimum of two matches per image is enforced to declare a correlated image pair, since this is the minimum that is required to unambiguously estimate a rotational offset between the pair.

In preparation for the global minimization step (see § 2.2 below), all point source match positions and uncertainties are mapped into the pixel coordinates of a Cartesian reference image frame. The definition of this reference depends on whether absolute or relative refinement is desired. The POINTINGREFINE software distinguishes between these options

⁴ The complete POINTINGREFINE package is available at <http://ssc.spitzer.caltech.edu/postbcd>.

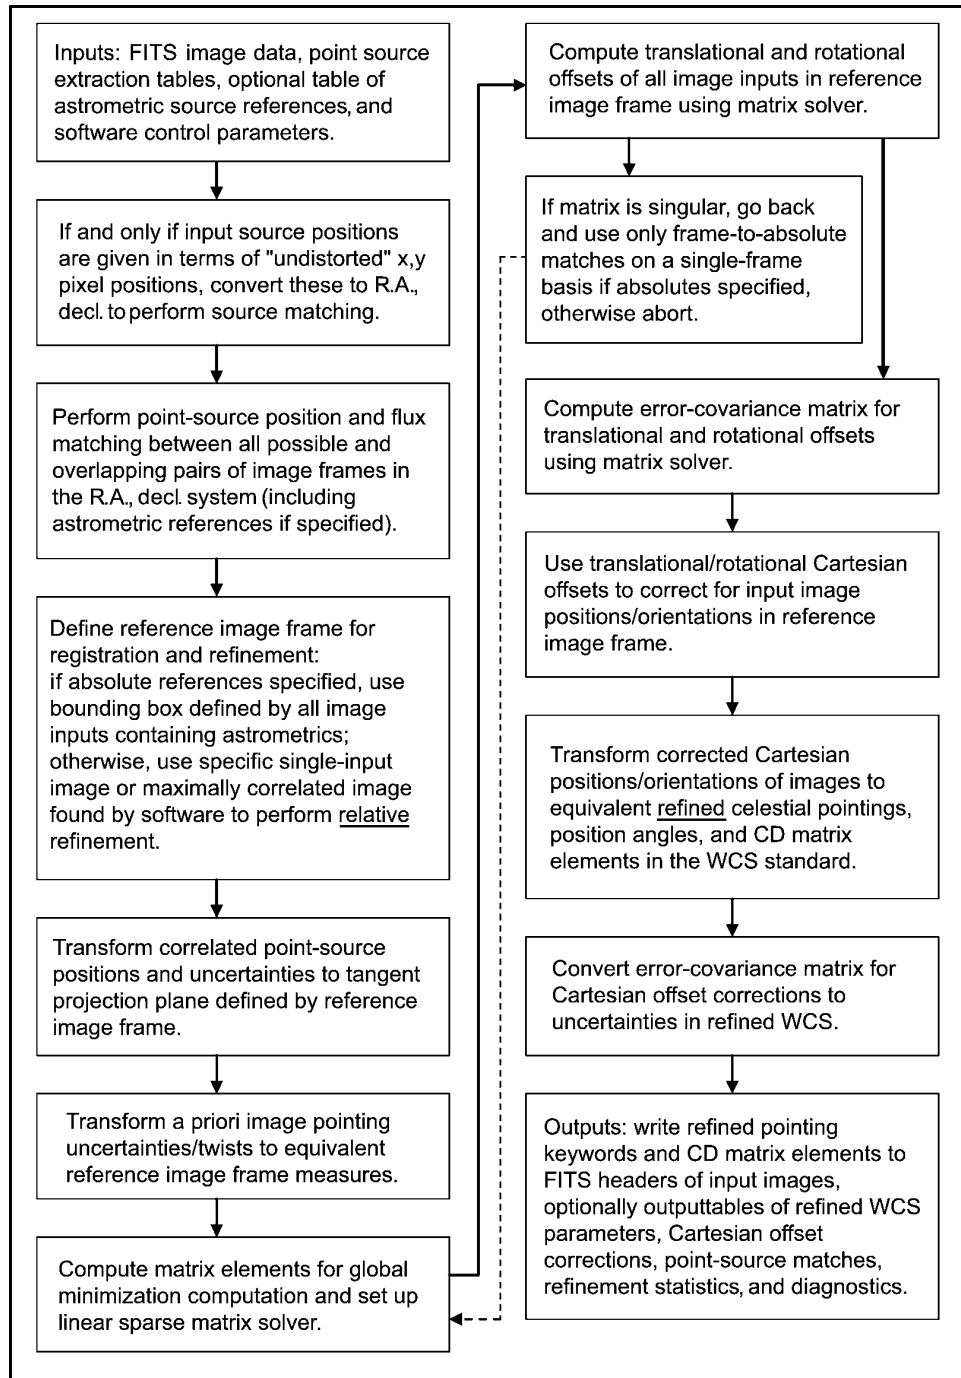


FIG. 1.—Processing and algorithmic flow in POINTINGREFINE software.




FIG. 2.—Simple three-image mosaic. Filled circles show sources detected from image n , and open circles show sources from m or 1.

from its given inputs. Two inputs are required for absolute refinement. First, the software expects a fiducial image frame (FIF). This is a file listing various WCS parameters that define the tangent point and boundary of a “fiducial image” encompassing all images. Second, a list of astrometric (absolute) sources that fall within the FIF boundary is required. The FIF acts as an effective input image whose sources are the astrometric references. Refinement with respect to such a FIF ensures both absolute and relative refinement among images. In relative refinement mode, all input images can be registered and refined with respect to a single input image. In this mode, the user has the option of making the software automatically select an image from the input list that is maximally correlated (has most overlaps) with other images, or they can specify their own.

2.2. Global Minimization

Consider the simple three-image mosaic in Figure 2. Image 1 defines the “fiducial” reference frame. The circles represent point sources detected from each overlapping image pair transformed into the reference frame of image 1. The filled circles are sources extracted from image n , and the open circles are sources extracted from either image 1 or m . The matched source pairs are shown offset from each other to mimic the presence of random pointing uncertainty in each input image. These are the offsets we wish to compute and use to correct each frame pointing. A one-dimensional representation of the projection geometry in the reference image frame is depicted in Figure 3. All input images have sizes in the reference frame that depend on their distance from the reference image tangent point. The projected linear size scales with angular distance θ , as $\sim 1 + \tan^2 \theta$. The POINTINGREFINE algorithm appropriately accounts for the inflation of centroid uncertainties and sepa-




FIG. 3.—One-dimensional representation of projection geometry showing images on sky and in tangent plane of reference image. Images 1 and 2 have the same physical size, but different projected sizes.

rations of correlated sources when projected into the reference frame. A potential problem for the algorithm is if image sizes and mosaic extents are large enough to cause a nonuniform dependence in scale over the region in which an input image is projected. The dependence of this effect with θ_i , however, is weak. For instance, the projected scale varies by $\leq 1\%$ across an image $30'$ in size at angular distances $\theta_i \leq 30^\circ$.

The main assumption of the algorithm is that the random uncertainty in the measured twist angle⁵ of an individual image frame ($\delta\theta^m$ in Fig. 2) is small enough to ensure $\sin \delta\theta^m \approx \delta\theta^m$ (see below). A twist angle uncertainty of $\delta\theta \leq 1'$ is a good working measure for the purposes of this algorithm. This is justified for the *Spitzer* science instruments, in which the absolute twist angle uncertainty (IPF Team 2003) is typically $\leq 30''$ (1σ).

In a rectilinear coordinate system (defined, say, by the frame of image 1 in Fig. 2), the positions of a point source i detected from images m and n are related via the transformation

$$\begin{pmatrix} x_i^m \\ y_i^m \end{pmatrix} \rightarrow \begin{pmatrix} \tilde{x}_i^m \\ \tilde{y}_i^m \end{pmatrix} = \begin{pmatrix} x_c^m \\ y_c^m \end{pmatrix} + \begin{pmatrix} \cos \delta\theta^m & -\sin \delta\theta^m \\ \sin \delta\theta^m & \cos \delta\theta^m \end{pmatrix} \begin{pmatrix} x_i^m - x_c^m \\ y_i^m - y_c^m \end{pmatrix} + \begin{pmatrix} \delta X^m \\ \delta Y^m \end{pmatrix}, \quad (2)$$

where $\delta\theta^m$ is a rotational offset, δX^m and δY^m are orthogonal translations, and (x_c^m, y_c^m) are the center coordinates of image m in the reference image frame. The rotation $\delta\theta^m$ is measured in a counterclockwise sense so that a rotation followed by orthogonal translations in x and y of image m will align the sources (*open circles*) detected therein with those detected in image n .

The assumption of small twist angle uncertainty (see above) implies $\sin \delta\theta^m \approx \delta\theta^m$ and $\cos \delta\theta^m \approx 1$, and thus the pair of

⁵ We define “twist angle” as the relative orientation of an image in an orthogonal coordinate system, *not* the conventional position angle measured in the ICRS.

equations defined by equation (2) can be linearized in $\delta\theta^m$ as

$$\begin{aligned}x_i^m &\rightarrow \tilde{x}_i^m = x_i^m - (y_i^m - y_c^m)\delta\theta^m + \delta X^m, \\y_i^m &\rightarrow \tilde{y}_i^m = y_i^m + (x_i^m - x_c^m)\delta\theta^m + \delta Y^m,\end{aligned}\quad (3)$$

where $(\tilde{x}_i^m, \tilde{y}_i^m)$ represents a position that has been corrected for unknown image rotation and orthogonal translations.

We define a cost function L , representing an inverse-variance weighted sum of squares of all matched point-source positional differences from all overlapping image pairs (m, n) :

$$\begin{aligned}L = \sum_{m, n(n>m)} \sum_i &\left[\frac{1}{\Delta x_i^{m, n}} (\tilde{x}_i^m - \tilde{x}_i^n)^2 \right. \\&\left. + \frac{1}{\Delta y_i^{m, n}} (\tilde{y}_i^m - \tilde{y}_i^n)^2 \right] + L_{\text{a priori}},\end{aligned}\quad (4)$$

where

$$\begin{aligned}\Delta x_i^{m, n} &= \sigma^2(x_i^m) + \sigma^2(x_i^n), \\ \Delta y_i^{m, n} &= \sigma^2(y_i^m) + \sigma^2(y_i^n),\end{aligned}\quad (5)$$

and σ^2 represent variances in extracted point-source centroids. The summations in equation (4) are over all matched point-sources from all image overlaps, including sources matched to any astrometric references present in the FIF, if available. The FIF, or absolute reference frame (as defined in § 2.1), is treated as an effective input image when source matching is performed.

The additive term $L_{\text{a priori}}$ in equation (4) represents an a priori weighting function that makes use of actual measured pointing uncertainties of the images. This function is defined as

$$\begin{aligned}L_{\text{a priori}} = \sum_{m, n(n>m)} &\left\{ \frac{(\delta X^m)^2}{\sigma_{Xm}^2} + \frac{(\delta Y^m)^2}{\sigma_{Ym}^2} + \frac{(\delta\theta^m)^2}{\sigma_{\theta m}^2} \right. \\&\left. + \frac{(\delta X^n)^2}{\sigma_{Xn}^2} + \frac{(\delta Y^n)^2}{\sigma_{Yn}^2} + \frac{(\delta\theta^n)^2}{\sigma_{\theta n}^2} \right\},\end{aligned}\quad (6)$$

where σ_{jm}^2 and σ_{jn}^2 ($j = X, Y, \theta$) represent *measured* pointing variances (in the ICRS) transformed into the reference image frame. The purpose of including $L_{\text{a priori}}$ is to avoid overrefining or biasing those images whose inherent measured pointing uncertainties are already small (within nominal requirements). In other words, those images whose uncertainties are known to be small a priori will have a larger contribution to $L_{\text{a priori}}$ relative to the correlated source term (double sum in eq. [4]). Consequently, the solution will be biased toward minimizing $L_{\text{a priori}}$ and not the correlated source term, which could potentially degrade expected image offsets (and final refined pointings). In the limit $L \rightarrow L_{\text{a priori}}$, the global minimum will be closer to $\delta\theta \approx \delta X \approx \delta Y \approx 0$. Conversely, large pointing uncertainties

will bias the solution toward the correlated source term in which refinement via point-source matches is obviously needed.

Equation (4) can be rewritten in terms of physical image offsets $\delta\theta^m$, δX^m , and δY^m for an arbitrary image m via equations (3) and (6):

$$\begin{aligned}L = \sum_{m, n(n>m)} &\left(\sum_i \left\{ \frac{1}{\Delta x_i^{m, n}} [x_i^m - (y_i^m - y_c^m)\delta\theta^m \right. \right. \\&\quad \left. \left. + \delta X^m - x_i^n + (y_i^n - y_c^n)\delta\theta^n - \delta X^n \right]^2 \right. \\&\quad \left. + \frac{1}{\Delta y_i^{m, n}} [y_i^m + (x_i^m - x_c^m)\delta\theta^m \right. \\&\quad \left. + \delta Y^m - y_i^n - (x_i^n - x_c^n)\delta\theta^n + \delta Y^n \right]^2 \Big\} \\&\quad \left. + \frac{(\delta X^m)^2}{\sigma_{Xm}^2} + \frac{(\delta Y^m)^2}{\sigma_{Ym}^2} + \frac{(\delta\theta^m)^2}{\sigma_{\theta m}^2} \right. \\&\quad \left. + \frac{(\delta X^n)^2}{\sigma_{Xn}^2} + \frac{(\delta Y^n)^2}{\sigma_{Yn}^2} + \frac{(\delta\theta^n)^2}{\sigma_{\theta n}^2} \right).\end{aligned}\quad (7)$$

The cost function defined by equation (7) can be treated as a standard χ^2 statistic to the extent that point-source centroid uncertainties are independently random and Gaussian, a good approximation in this case. The probability density function for L about its global minimum is therefore the χ^2 distribution for ν degrees of freedom (dof) where

$$\nu = 2N_{\text{matches}} - 3(N_{\text{imgs}} - 1).\quad (8)$$

N_{matches} is the total number of point-source matches in all image overlap regions, and N_{imgs} is the total number of images containing the detected matches (including the reference image). The total number of parameters to solve is actually $3(N_{\text{imgs}} - 1)$, since we have three offsets ($\delta\theta^m$, δX^m , and δY^m) for every image m , and we exclude the reference image, which by definition is constrained to have $\delta\theta = 0$, $\delta X = 0$ and $\delta Y = 0$.

Our aim is to minimize L with respect to all image offsets $\delta\theta^m$, δX^m , and δY^m for every correlated image m . At the global minimum of L , partial derivatives with respect to these three offsets for each image m are required to vanish:

$$\frac{\partial L}{\partial \delta\theta^m} = 0; \quad \frac{\partial L}{\partial \delta X^m} = 0; \quad \frac{\partial L}{\partial \delta Y^m} = 0.\quad (9)$$

Evaluating these partial derivatives leads to a set of three simultaneous equations for each image in the ensemble. For N_{imgs} , we therefore have $3(N_{\text{imgs}} - 1)$ simultaneous equations in $3(N_{\text{imgs}} - 1)$ unknowns. Each image m of an ensemble $m = 1 \dots M$ could potentially be correlated (contain sources in common) with any other image n in which $n =$

$1 \dots N; m \neq n$ and $M = N = N_{\text{imgs}} - 1$. From the conditions defined in equation (9), the simultaneous system of equations used to solve for the offsets of all images m can be represented as the matrix equation

$$\begin{pmatrix} A_{\theta}^{m=1} & A_X^{m=1} & A_Y^{m=1} & \dots & A_{\theta}^{m=N} & A_X^{m=N} & A_Y^{m=N} \\ B_{\theta}^{m=1} & B_X^{m=1} & 0 & \dots & B_{\theta}^{m=N} & B_X^{m=N} & 0 \\ C_{\theta}^{m=1} & 0 & C_Y^{m=1} & \dots & C_{\theta}^{m=N} & 0 & C_Y^{m=N} \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ A_{\theta}^{m=M} & A_X^{m=M} & A_Y^{m=M} & \dots & A_{\theta}^{m=M} & A_X^{m=M} & A_Y^{m=M} \\ B_{\theta}^{m=M} & B_X^{m=M} & 0 & \dots & B_{\theta}^{m=M} & B_X^{m=M} & 0 \\ C_{\theta}^{m=M} & 0 & C_Y^{m=M} & \dots & C_{\theta}^{m=M} & 0 & C_Y^{m=M} \end{pmatrix} \times \begin{pmatrix} \delta\theta^{m=1} \\ \delta X^{m=1} \\ \delta Y^{m=1} \\ \dots \\ \delta\theta^{m=M} \\ \delta X^{m=M} \\ \delta Y^{m=M} \end{pmatrix} = \begin{pmatrix} \Psi_A^{m=1} \\ \Psi_B^{m=1} \\ \Psi_C^{m=1} \\ \dots \\ \Psi_A^{m=M} \\ \Psi_B^{m=M} \\ \Psi_C^{m=M} \end{pmatrix}. \quad (10)$$

Equation (10) is of the form $\mathbf{M} \cdot \mathbf{X} = \Psi$. Elements of the coefficient matrix \mathbf{M} and the right-hand side column vector Ψ are given in Appendix A. The solution for the column vector \mathbf{X} with $3(N_{\text{imgs}} - 1)$ unknowns can be obtained using linear sparse matrix methods, since a large number of the matrix elements can be zero, depending on the mosaic geometry. In general, the fraction of zeros will be $\geq \frac{2}{9}$ ($\geq 22\%$). The $\sim 22\%$ minimum will occur when every image of the ensemble is correlated with every other image, such as in a co-added stack. The level of sparsity in \mathbf{M} will increase with nonzero elements along a block diagonal if one desires to tie and refine images with respect to astrometric absolute references *alone*. In this specialized case, $n =$ the reference image, and all elements with superscript n in \mathbf{M} will be zero. In general, the minimum and maximum possible number of zeros in the matrix \mathbf{M} are

$$\begin{aligned} N_{\min} &= 2(N_{\text{imgs}} - 1)^2, \\ N_{\max} &= 9(N_{\text{imgs}} - 1)^2 - 7(N_{\text{imgs}} - 1). \end{aligned} \quad (11)$$

We use the UMFPAK⁶ library to solve the matrix equation (10). This is adapted for solving large nonsymmetric matrix systems. The library includes an iterative scheme to correct solutions for the inevitable accumulation in round-off error during the Lower and Upper (LU) matrix factorization stage.

We also compute the full error-covariance matrix for all image offsets ($\delta\theta^m$, δX^m , and δY^m), which is given by the inverse of the coefficient matrix, $\mathbf{C} = \mathbf{M}^{-1}$ (e.g., Press et al. 1999). Variances in each offset are along the diagonal of \mathbf{C} , and covariances are given by off-diagonal elements. The covariance matrix \mathbf{C} is computed using the same matrix solver as above

⁶ See <http://www.cise.ufl.edu/research/sparse/umfpack>.

by repeatedly solving for each unknown column \mathbf{X}_c in \mathbf{M}^{-1} such that $\mathbf{M} \cdot \mathbf{X}_c = \mathbf{I}_c$, where \mathbf{I}_c is the corresponding column in the identity matrix. An analysis of the covariance matrix using real data is presented in Appendix B.

Images that are *not* correlated (lack point-source matches) with any others in the input ensemble cannot contribute to the globally minimized cost function L . For these images, the POINTINGREFINE software explicitly sets their reference frame offsets to zero, and no refinement of their pointing is possible. As a further detail, the inclusion of the a priori weighting function $L_{\text{a priori}}$ (eq. [6]) guarantees that the matrix in equation (10) will be nonsingular. The priors provide at least one measurement per image. If these are omitted from the cost function (eq. [4]), then there are cases in which the matrix could be singular. This can occur if the input image ensemble contains clusters of correlated images disjoint from each other and with a noncontiguous/broken path existing between the clusters. This situation leads to an underrepresentation of images across the full simultaneous system of equations, and the determinant will be zero. As indicated in the processing flow of Figure 1, if this occurs in “absolute refinement” mode, a second-pass computation is attempted, and only those frames that contain absolute astrometric matches are used. No relative frame-to-frame information is used, and images are refined based on their absolute source content alone. The matrix effectively becomes block diagonal, and a nonzero determinant is guaranteed. In “relative refinement” mode, no attempt is made to perform registration within each disjoint subensemble. Instead, the software will abort with a message indicating that disjoint clusters exist.

Once image offset corrections ($\delta\theta^m$, δX^m , and δY^m) and associated uncertainties have been determined in the reference image frame, the final step involves refining the celestial pointing and orientation of each image m . This is performed by correcting the pointing centers (x_c^m , y_c^m) of each image in the reference image frame via equation (3); i.e.,

$$\begin{aligned} x_c^m(\text{new}) &= x_c^m(\text{old}) + \delta X^m, \\ y_c^m(\text{new}) &= y_c^m(\text{old}) + \delta Y^m, \end{aligned} \quad (12)$$

and then transforming back to the sky to yield refined pointings in the ICRS. Image orientations are refined in a similar manner. For these, we correct and transform two fiducial points per image to uniquely determine the refined position angle. The main outputs of the POINTINGREFINE software are additional WCS keywords written to FITS image headers representing refined pointings and orientations on the sky (see processing flow in Fig. 1).

2.3. Optimization and Expectations

The accuracy in pointing refinement or registration can be severely limited by systematics. For example, inaccurately calibrated image scale and/or distortions known a priori to be po-

sition dependent will bias extraction centroids and have adverse effects on the separations of bona fide source matches and, hence, final globally minimized solutions. A method to test for possible contamination from systematics involves examining distributions of matched source separations after refinement and comparing them with those contributed by (random) centroid uncertainties. This is performed with real data in § 4.

The presence of absolute-astrometric point sources are an important ingredient for refinement and registration in general. These reduce the potential for a biased random walk in the refined pointing of an image as its distance from a fiducial reference increases. Such effects are also alleviated by using appropriate prior pointing uncertainty information in the global minimization cost function (see eq. [4]). Absolute astrometric references provide a set of “anchor points” to which all extractions will be attracted. If one wants to perform *absolute* refinement with few or no astrometric references, then sufficient numbers of images are needed to increase the number of potential matches in frame overlap regions for a single point source. This is necessary in order to approach a normal distribution about the expected absolute source position. In other words, large numbers of correlated source positions will ensure that the mean position of a correlated source cluster is close to the “truth,” or that which will result after refinement (assuming there are no systematics as discussed above). Because of the rarity of cases involving numerous image overlaps to provide good normally distributed sampling, it is advisable to use astrometric references wherever possible.

The accuracy to which we can refine the pointings of an ensemble of mutually correlated images primarily depends on the number of point-source matches available (both relative and absolute). With randomly distributed uncertainties in point-source centroids, it is expected that the mean separation between matched source positions is approximately Gaussian after refinement, by virtue of the central limit theorem. In this limit, the (radial) uncertainty in refinement⁷ of a single image will scale as

$$\sigma_r \approx \sqrt{\frac{\sigma_{\text{ext}}^2 + \sigma_{\text{abs}}^2}{(N_{\text{ext}}/2)[1 + (\sigma_{\text{abs}}^2/\sigma_{\text{ext}}^2)] + N_{\text{abs}}}}, \quad (13)$$

where N_{ext} and N_{abs} respectively represent the number of frame-to-frame and frame-to-absolute source match pairs in *all* overlap regions associated with the image, σ_{ext} is a typical source extraction centroid uncertainty, and σ_{abs} is an astrometric source position uncertainty. This approximation assumes that $\sigma_{\text{ext}} > 0$ or $\sigma_{\text{abs}} > 0$ when either $N_{\text{ext}} > 0$ or $N_{\text{abs}} > 0$, respectively. For cases in which there are no astrometric matches, we set $N_{\text{abs}} = 0$ and $\sigma_{\text{abs}} = 0$, and equation (13) reduces to $\sigma_{\text{ext}}(2/N_{\text{ext}})^{1/2}$. We expect

⁷ Approximated as the uncertainty in the mean source match separation, with source positions weighted by their inverse variances.

to measure source extraction centroids to better than ~ 0.1 pixel, ($\sim 0''.121$ for *Spitzer*'s IRAC focal plane arrays). If we assume, for example, astrometric positional errors of $\sim 0''.2$ (conservatively speaking), then to refine image pointings to an accuracy better than $\sim 0''.1$ will require at least five *astrometric* point-source matches per frame if $N_{\text{ext}} = 0$, or less if $N_{\text{ext}} > 0$. One can see that the refinement accuracy increases with more point-source matches. This assumes that the observational setup allows for sufficient frame-to-frame overlap to ensure good numbers of relative matches N_{ext} . If this is not the case, one will have to resort to using pure astrometric (absolute source) matches alone.

To summarize, corrections for optimal refinement will effectively be given by the *magnitude* of frame-pointing uncertainties, with errors approximated by equation (13). The latter assumes that in the limit of the increasing number of matches, the mean source separation per image overlap region is approximately normally distributed, with an uncertainty determined exclusively by point-source centroid uncertainties. Any position-dependent systematic offset between source matches, such as nonuniform pixel scale or inaccurately calibrated distortion, will limit the refinement accuracy to the size of systematic error involved.

3. VALIDATION USING A MONTE CARLO SIMULATION

We quantitatively assess the performance of the above algorithm by using a simulation of 1000 mosaicked images, with each image's coordinates modeled with an uncertainty drawn from a Gaussian distribution. The simulation is generic in the sense that it represents a good overall representation of the type of data that could be acquired with modern optical/near-infrared detectors to moderately faint magnitudes ($m_{\text{opt}} \approx 23$ or $m_{\text{near-IR}} \approx 19$). To facilitate a comparison with real observations in § 4, we have decided to model the source count distribution and detector properties with that expected (and more or less observed) in the $3.6 \mu\text{m}$ band of *Spitzer*'s IRAC instrument. A more detailed description is given in § 4.

The “truth” source flux-density distribution was simulated using models given by Xu et al. (2003). These assume a high galactic latitude stellar model, several galaxy luminosity functions that depend on galaxy morphological type, and exploit a large library of spectral energy distributions. This simulation was used extensively for predicting *Spitzer* source populations (Lonsdale et al. 2003) and for validating processing pipelines.

The simulation steps are as follows.

1. A “truth” list of random source flux densities was generated using models from Xu et al. (2003), covering an area of 0.45 deg^2 down to a flux density of $10 \mu\text{Jy}$ at $3.6 \mu\text{m}$ (equivalent to $\approx 18.6 \text{ mag}$ in the Vega system).

2. Truth sources were assigned both random and correlated positions within the area to be mosaicked. Galaxies were assumed to have a weak correlated component, with amplitude (excess above random) based on an empirical *K*-band two-point

angular correlation function $w(\theta) \approx 0.001(\theta/\text{deg})^{-0.8}$. A pixel scale of $\sim 1''.21$ (characteristic of the IRAC arrays) was used when mapping sources to the pixel frame of the mosaic.

3. Truth sources were convolved with a PSF⁸ that was scaled by the appropriate source flux. PSF-convolved truth sources were mapped into the mosaic frame, with no pixel resampling.

4. A total of 1000 (256×256 pixel) “truth image” frames were generated from the full mosaic area by using a dither and mapping strategy that assumed $\sim 60\%$ adjacent image overlap. The optimal map geometry for this 1000 image set was generated by D. Shupe (2002, SSC, private communication). R.A., decl., and twist angle information is also attached to each image at this stage.

5. A new set of 1000 images was generated (our simulated control sample), with uncertainties added to image pointings and position angles. Pointing offsets (prior uncertainties) were modeled as Gaussian random deviates along each independent orthogonal image axis. These were drawn from a zero-mean Gaussian distribution with $\sigma = 0''.85$ per axis. This choice for σ is based on a prelaunch pointing knowledge of $\sim 1''.2$ (1σ radial) for *Spitzer* in the science instrument frame (Lawrence et al. 2000). Even though *Spitzer* can now actually do better than this by a factor of ~ 1.5 (see § 4), the ultimately refined pointings are independent of the magnitude of any reasonable simulated prior uncertainty assumed (see below). Twist angle uncertainties were modeled as Gaussian, with $\sigma = 20''$. A smoothly varying background adjusted with the expected Poisson and (IRAC band 1) read noise per pixel was added to each image.

6. An absolute source list (representing astrometric references) was generated by taking the brightest true sources that gave a mean density of ≈ 50 sources per $5''.2 \times 5''.2$ image region. This resulted in 3030 true sources. To imitate an astrometric catalog, the sources were assigned positions modeled as truth $\pm u\sigma$ along each axis, with u drawn from a Gaussian distribution with $\sigma = 0''.06$. This is typical for sources in the Two Micron All Sky Survey (2MASS) Point-Source Catalog to $K_s \approx 15$ (see § 4 for details regarding the 2MASS catalog).

7. The SSC point-source extractor was used on each simulated (control) image to extract sources above a threshold of 5σ . This resulted in ~ 40 extractions per frame.

The POINTINGREFINE software was executed on the 1000 image control sample. A source match radius of $3''.5$ was used to comfortably accommodate prior image pointing errors and extraction centroid errors (typically $0''.15$, 1σ per axis). Simultaneous flux matching was also applied between frame-to-frame and frame-to-absolute (astrometric) matches, with maximum flux-difference thresholds of 5% and 10%, respectively. A zoomed-in ($2''.8 \times 3''.8$) section of our 1000 image simulation (with image pointing errors) is shown in the left panel of Figure




FIG. 4.—*Top*: Unrefined mosaic section of 1000 image IRAC simulation at $3.6\ \mu\text{m}$; *bottom*: same section after pointing refinement. Field dimensions are $\approx 2''.8 \times 3''.8$.

4. On the right is the same section after pointing refinement. The increase in resolution is dramatic. There is a factor of ~ 6 decrease in mean source match separation, leading to more localized point-source flux distributions and detectability to fainter levels. In this test, the surface brightness is increased by factors of ~ 2 – 3.5 for detected sources after refinement. This is as expected, given that the 1σ radial image pointing uncertainty is of the order the input pixel size, and the intrinsic PSF has a FWHM of $\approx 1''.66$.

To get a more quantitative assessment of the performance of POINTINGREFINE, we compare the distribution of separations between image center pointings of “truth” and simulated (control sample) images before and after refinement. This is shown in the top panel of Figure 5. Two different runs of POINTINGREFINE were performed, based on the number of (brightest) extractions used per image. One gave an average of ~ 10 relative (frame to frame) and ~ 20 absolute-astrometric source matches per image (the “10/20 match” case), and the second resulted in an average of ~ 2 relative and ~ 3 absolute matches per image (the “2/3 match” case). It should be noted that two matches per image is the absolute minimum required to unambiguously determine two orthogonal shifts and an orientation per frame (i.e., the number of degrees of freedom

⁸ Made from in-flight IRAC band 1 observations by the IRAC Instrument Support Team at the SSC. It has an FWHM of $\sim 1''.66$ and $\sim 42\%$ central pixel flux.




FIG. 5.—*Left*: Distributions of image center separations relative to “truth” before and after refinement. *Right*: Offsets in R.A. and decl. between image centers relative to truth. Circle represents 2σ region of “truth refined” distribution for case with ~ 2 relative/3 absolute matches per image (dot-filled histogram in top figure).

(dof) = $2N_{\text{matches}} - 3 = 1$). The dispersion in image separation relative to truth *after refinement* for the 10/20 match case is ≈ 65 mas (1σ radial). For the 2/3 match case, this is ≈ 110 mas. For the given match statistics, these numbers are more or less consistent with the simple scaling relation given by equation (13). Offset distributions along each axis are shown in the lower panel of Figure 5, where the open circle represents the 2σ contour for the 2/3 match case.

By comparing differences in image radial separations before and after refinement, we find that over 890 images have their pointing refined to better than 95% relative to “truth” (for the 10/

20 match case). The top panel in Figure 6 shows the distribution in fractional refinement. This quantity is defined as the ratio of separations, $1 - D(\text{refined} - \text{truth})/D(\text{unrefined} - \text{refined})$. For the 2/3 match case, slightly less than half have the same amount of refinement, although most images are refined to better than 80%, corresponding to a discrepancy of ≈ 180 mas within truth image positions. Thus, the refinement is very good, even with minimal matches.

The bottom panel in Figure 6 shows distributions in matched source radial separations before and after refinement. Uncertainties in both image pointing and source extraction centroids

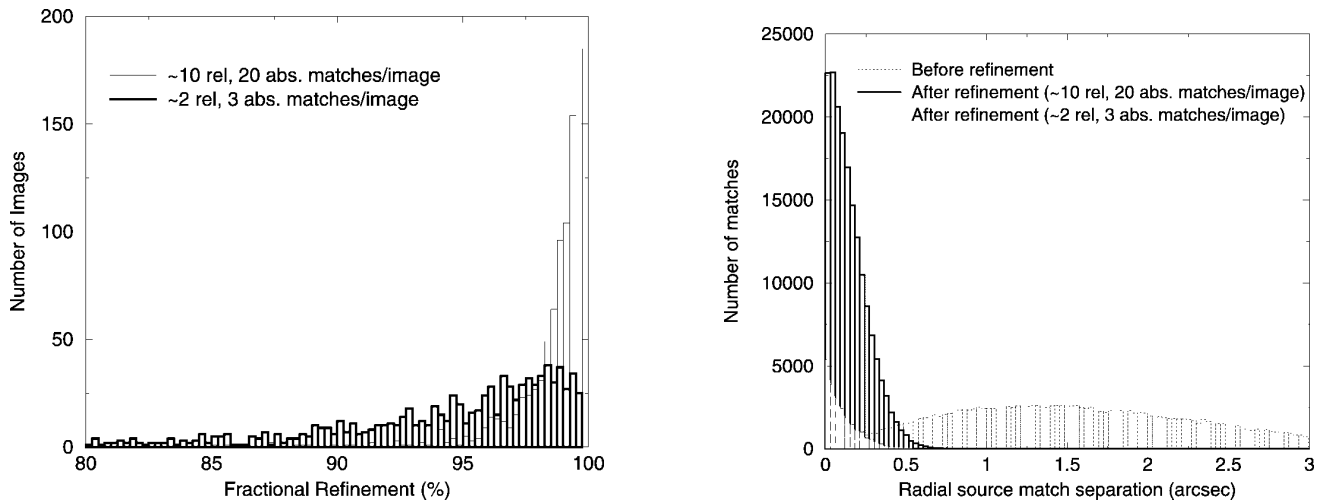


FIG. 6.—*Left*: Distributions of the magnitude of refinement, represented as a percentage of initial “truth unrefined” image separation. *Right*: Distributions of matched source radial separations before and after refinement.

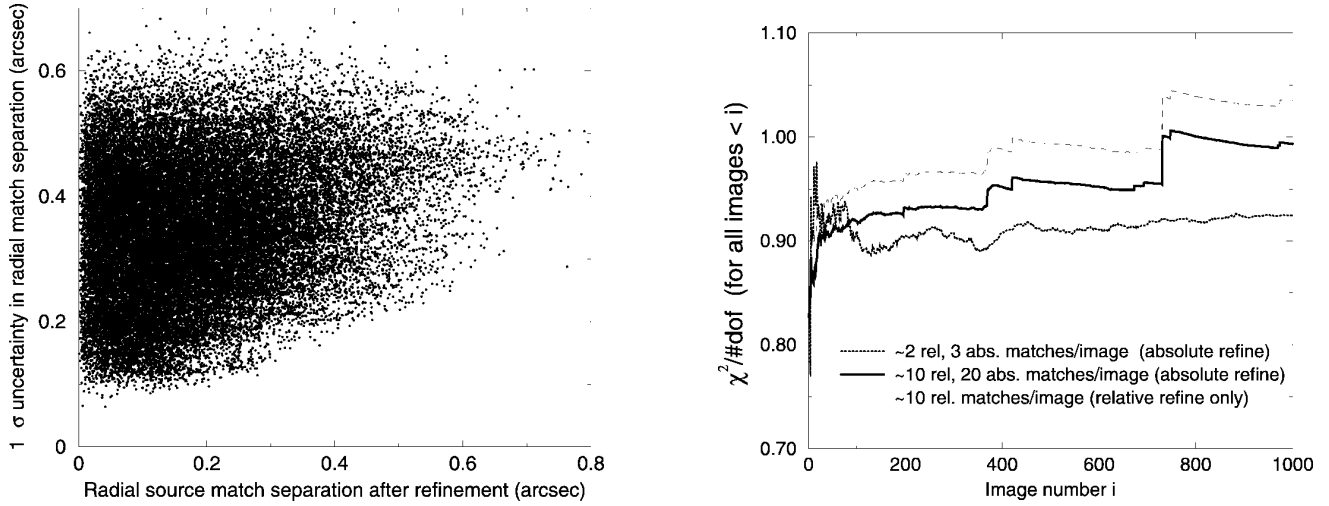


FIG. 7.—*Left*: Uncertainty in matched source radial separation (from centroid uncertainties) as a function of actual source separation after refinement. *Right*: Reduced χ^2 (i.e., χ^2/dof) as a function of mosaic subset composed of successively increasing numbers of images i (see § 3).

(or radial separation) contribute to the dispersion in the *unrefined* distribution. Image pointing uncertainties dominate, with a $1''/2$ contribution compared to $\approx 0''/2$ for extraction centroids (both 1σ radial). After pointing refinement, the dispersion in radial source separation is expected to be dominated exclusively by extraction centroid errors and, indeed, the distributions confirm this for both the $2/3$ and $10/20$ match cases (Fig. 6, *bottom left*). The distribution for the $2/3$ match case is narrower, since such matches were performed using the brightest extractions per image, and these inherently have better determined centroids.

The separation between two source extraction centroids (x_i, y_i) and (x_j, y_j) is given by

$$r = [(x_i - x_j)^2 + (y_i - y_j)^2]^{1/2}. \quad (14)$$

Since a majority of extractions are unresolved point sources with very circular error ellipses, we can ignore any correlations between uncertainties along each axis. To a good approximation, the uncertainty in radial separation between any two centroids can therefore be written

$$\sigma_r \approx \sqrt{\sigma_{x_i}^2 + \sigma_{x_j}^2} \approx \sqrt{\sigma_{y_i}^2 + \sigma_{y_j}^2}, \quad (15)$$

where $(\sigma_{x_i}^2, \sigma_{y_i}^2)$ and $(\sigma_{x_j}^2, \sigma_{y_j}^2)$ are centroid variances in each axis for sources i and j , respectively. A comparison between uncertainties in matched source radial separation (eq. [15]) and actual separations (eq. [14]) after refinement is shown in the top panel of Figure 7 (for $10/20$ match case). They are both mutually consistent, although the spread is greater at separations $\lesssim 0''/3$. After refinement and in the absence of systematics, any residual separation in a matched source pair must be due to intrinsic centroiding error alone. Since separations between

matched sources along each axis are, to a good approximation, independently random and normally distributed with zero mean ($\langle x_i - x_j \rangle = \langle y_i - y_j \rangle = 0$), the quantity r (eq. [14]) can be shown to follow a χ distribution with 2 dof (e.g., Evans et al. 2000). This special case is also known as the Rayleigh distribution:

$$P(r) = \frac{r}{\beta^2} \exp\left[-\frac{1}{2}\left(\frac{r}{\beta}\right)^2\right], \quad (16)$$

where β is a parameter characterizing the width. This can be written in terms of the second moment (variance) of $P(r)$ as

$$\beta = \sqrt{\left(\frac{2}{4 - \pi}\right) \sigma_r^2} \approx 1.52\sigma_r. \quad (17)$$

For a given uncertainty σ_r as computed from equation (15), the variation in the density of points with r along any horizontal cut in Figure 7 is qualitatively consistent with that predicted by equations (16) and (17).

The bottom panel of Figure 7 shows the dependence of the reduced χ^2 (effectively the cost function in eq. [7] divided by number of degrees of freedom defined by eq. [8]) as a function of increasing number of images N_i in our 1000 image simulation. It is important to note that image offsets are not recomputed using repeated global minimizations for each new set of images N_i . Instead, the original full 1000 image solution of image offsets is used throughout to recompute χ^2 from equation (7) as N_i is increased for all image pairs (m, n) such that $n < m \leq N_i$. As one approaches the full image set of $N_i = 1000$, one expects the reduced χ^2 to converge to unity if, on average, residuals in source separations after refinement are purely consistent with extraction centroid uncertainties. The lower re-




FIG. 8.—Coverage mosaic and mapping geometry of the 105 image IRAC observation used in this analysis. Adjacent images have $\sim 20\%$ overlap, with a coverage of six and 12 pixels at edges and (inner) corners, respectively. Mapping direction is shown by arrows starting at top right and ending at bottom left. Open circles are image centers.

duced χ^2 values for smaller image numbers (and particularly for all image numbers in the 2/3 match case) is due to the nonlinear behavior in χ^2 when the number of degrees of freedom is small. Better fits (smaller χ^2) can be obtained for very low numbers of degrees of freedom. In Figure 7 (*bottom panel*), only the 10/20 match case (with absolute refinement as studied above) shows approximate convergence to 1, and the other curves are not far from it. Also shown is a case in which only relative frame-to-frame matches and no astrometric references are used. This case tends to show a slightly higher reduced χ^2 (≈ 1.04), which is significant, since it is almost 20 standard deviations from the expected value in reduced χ^2 ($\sigma_{\chi^2} = (2/\text{dof})^{1/2}$, where $\text{dof} = 308,935$). This was traced as being due to slightly underestimated uncertainties in extraction centroids. This is not seen in the absolute refinement case (*solid curve*), since absolute astrometric uncertainties are themselves overestimated, and their (almost equal) contribution tends to lower the effective χ^2 when combined with relative frame-to-frame matches.

To summarize, we have presented a simulation to ascertain the performance of the POINTINGREFINE algorithm. The model-dependent parameters entering our simulation can be

isolated to properties of $3.6 \mu\text{m}$ source populations, specifics of the IRAC band 1 array, such as PSF and pixelization, and a priori telescope pointing knowledge. These can be appropriately rescaled to model other wavelengths and detectors. However, in the absence of systematics in the locations of potential matches between frames (both absolute and relative), and regardless of instrumental setup or detector properties, the accuracy in refined pointing is purely dictated by the accuracy of point-source centroids and match statistics. Our simulation (Figs. 5 and 6) indeed shows that the refined pointing will typically never be worse than the (combined) centroid uncertainties of matched pairs of sources. A well-sampled and characterized PSF is expected to give centroiding accuracies to better than 1/10 of a resolution element. If errors are independently random, good match statistics can then only work in our favor to give the desired $\sim 1/(N_{\text{matches}})^{1/2}$ improvement in pointing accuracy.

4. TESTING ON SPITZER IRAC DATA

The Infrared Array Camera is one of three focal plane instruments on the *Spitzer Space Telescope* (Fazio 2004). IRAC provides simultaneous $\sim 5'2 \times 5'2$ images at 3.6, 4.5, 5.8, and $8 \mu\text{m}$ (bands 1–4). All four detector arrays in the camera are 256×256 pixels in size, with a pixel size $\approx 1''.2$. We present here the results of a case study of observations acquired with IRAC during the in-orbit checkout period (2003 October). In this section, we validate the pointing performance of IRAC and estimate the accuracy of refinement that can be achieved using a standard astrometric catalog and comparisons with our simulation of the $3.6 \mu\text{m}$ band from § 3.

The observational request used for our case study consists of 105 regularly spaced images arranged in a rectangular raster pattern. The coverage map and geometry is shown in Figure 8. Adjacent images have $\sim 20\%$ overlap in each axis, with a coverage of six and 12 pixels at the edges and (inner) corners, respectively. The mapping was performed in repetitive horizontal scans, as shown by arrows in Figure 8. The first image is at top right, and the last at bottom left. All images across all bands were preprocessed for removal of instrumental signatures using the SSC's IRAC pipeline,⁹ and raw pointing information was attached to FITS headers. Source extraction was then performed using the SSC source extractor and PSFs characterized from in-flight data. Sources were extracted to a uniform S/N of 5σ in each band, resulting in an average of $\sim 39, 24, 13,$ and 6 extractions per image for bands 1, 2, 3, and 4, respectively. Errors in flux-weighted centroids were on average $\sim 0''.18, 0''.22, 0''.26,$ and $0''.27$ (1σ radial) for each band, respectively.

We used data from the Two Micron All Sky Survey¹⁰ to define a standard astrometric catalog for all IRAC bands. Zacharias et al. (2003) used the USNO CCD Astrograph Catalog (UCAC; which is accurate to ~ 20 mas) to perform an assess-

⁹ See <http://ssc.spitzer.caltech.edu/documents/SOM>.

¹⁰ See <http://www.ipac.caltech.edu/2mass>.

TABLE 1
STATISTICS FOR IRAC OBSERVATION CASE STUDY

λ Band (μm)	$\langle\text{Abs.}\rangle^a$	$\langle\text{Rel.}\rangle^b$	$\langle D_{\text{bet}A} \rangle^c$ (arcsec)	$\langle D_{\text{bet}R} \rangle^c$ (arcsec)	$\langle D_{\text{at}T} \rangle^d$ (arcsec)	χ^2	dof ^e	χ^2/dof
3.6	34.7	20.5	0.847	0.605	0.151	8580.84	10821	0.792
4.5	20.7	13.7	0.883	0.622	0.152	3843.93	5645	0.681
5.8	9.4	7.6	0.933	0.603	0.142	1250.31	2443	0.511
8.0	5.3	2.1	0.913	0.651	0.203	1240.50	1249	0.993

^a Average number of absolute matches per frame.

^b Average number of relative (frame to frame) matches per frame.

^c $\langle D_{\text{bet}A} \rangle$ and $\langle D_{\text{bet}R} \rangle$ represent the mean matched source separation *before* refinement for absolute and relative matches, respectively (see Fig. 9).

^d $\langle D_{\text{at}T} \rangle$ represents the mean total matched source separation (for *both* relative and absolute matches) *after* refinement.

^e Number of degrees of freedom (see eq. [8]).

ment of the accuracy of 2MASS astrometry. They found random errors of 2MASS positions of $\sim 85\text{--}140$ mas radial to a limiting K_s magnitude of 15. We use astrometric sources detected in the 2MASS K_s band ($2.12 \mu\text{m}$), which is well suited to IRAC, and as a good working measure, we impose a magnitude limit of $K_s = 15$. To this limit and at the highest galactic latitudes ($|b| > 60^\circ$), one expects to find at least two astrometric sources per $\sim 5.2 \times 5.2$ IRAC field $\sim 60\%$ of the time. In the galactic plane anticenter, this increases to at least 30 sources 80% of the time (J. Surace 2000, SSC, private communication).

Using source extractions and astrometric references, the POINTINGREFINE software was executed on each of the four band-dependent 105 image ensembles. Guided by an ex post facto analysis of typical source separations and flux differences, point-source matching was performed using a nominal search radius of $2''$, and simultaneous flux matching was performed with relative (frame to frame) and astrometric (frame to absolute) flux difference thresholds of 4% and 50%, respectively, across all bands. Flux matching thresholds were set conservatively high to account for relative and absolute photometric calibration errors and intrinsic scatter between source populations detected in the 2MASS and IRAC bands. Relative and astrometric match statistics are summarized in Table 1.

Figure 9 shows distributions in matched source radial separations before and after refinement for all bands. Relative and astrometric matches have been separated. These distributions provide a powerful diagnostic with which to assess the in-flight pointing performance in the IRAC science instrument frame. The end-to-end pointing accuracy is a function of the inherent star-tracker accuracy, the spacecraft control system, how well the star-tracker bore sight is known in the telescope pointing frame and focal plane array (science instrument) frame, and variations in these due to thermomechanical deflections. *Spitzer's* star-tracker assembly alone provides pointing and control to better than $0''.3$ absolute accuracy over a 200 s integration (Lawrence et al. 2000). For comparison, each pointed observation in this study corresponds to an integration of ≈ 1.2 s.

We can get an estimate of both the relative and absolute raw pointing for IRAC from the “before refinement” (*heavy solid*

and dotted lines) histograms in Figure 9. The main difference between bands is in the number of point-source matches. Other than that, ranges in radial distributions are more or less consistent. Across all bands, the 1σ radial separation between frame-to-absolute matches is $\sim 0''.85\text{--}0''.93$, while for relative (frame to frame) matches, this is $\sim 0''.61\text{--}0''.65$ (See Table 1). In addition to actual pointing dispersion, these estimates include a dispersion component from point-source centroid errors. The contribution from centroiding error to the “before refinement” distributions, however, is negligible. Taking for instance the band 1 average centroiding error of $\sim 0''.18$ and $\sim 0''.085$ for 2MASS (1σ radial), this translates to effective match separation uncertainties of $\sim 0''.14$ and $\sim 0''.18$ for frame-to-absolute and frame-to-frame (relative) matches, respectively, assuming errors in each axis are *uncorrelated* (see eq. [15]). This implies that the actual absolute and relative pointing of the IRAC arrays is typically $(0.85^2 - 0.14^2)^{1/2} \approx 0''.83$ and $(0.61^2 - 0.18^2)^{1/2} \approx 0''.58$ (1σ radial), respectively. Similar results are found using other bands. These estimates are consistent or better than actual uncertainties reported by SSC pointing reconstruction software ($\approx 1''.1$ absolute in each axis). These account for uncertainty in the attitude of the telescope boresight system, which includes jitter due to limit cycle motions and uncertainty in the location of an off-axis science array due to uncertainty in boresight twist angle (both $\leq 0''.04$ rms), time-dependent biases between the telescope and star-tracker systems ($\geq 0''.6$ rms over 200 s), and the relationship between boresight to science array pointing frame ($\sim 0''.1$ rms) (Lawrence et al. 2000).

The source separation distributions *after* refinement in Figure 9 allow us to validate how well (or whether) images were refined to within accuracies determined by source extraction centroids. As found in our simulation (Fig. 5), images were refined to better than ≈ 65 mas (1σ radial) within truth pointings, with a corresponding ≈ 200 mas dispersion in source separations after refinement (Fig. 6). The top panel in Figure 10 shows the distribution of uncertainties in radial match separation (from eq. [15]) as a function of separation r for band 1 *after refinement*. These are broadly consistent above the minimum cutoff uncertainty of ~ 100 mas imposed by the finite

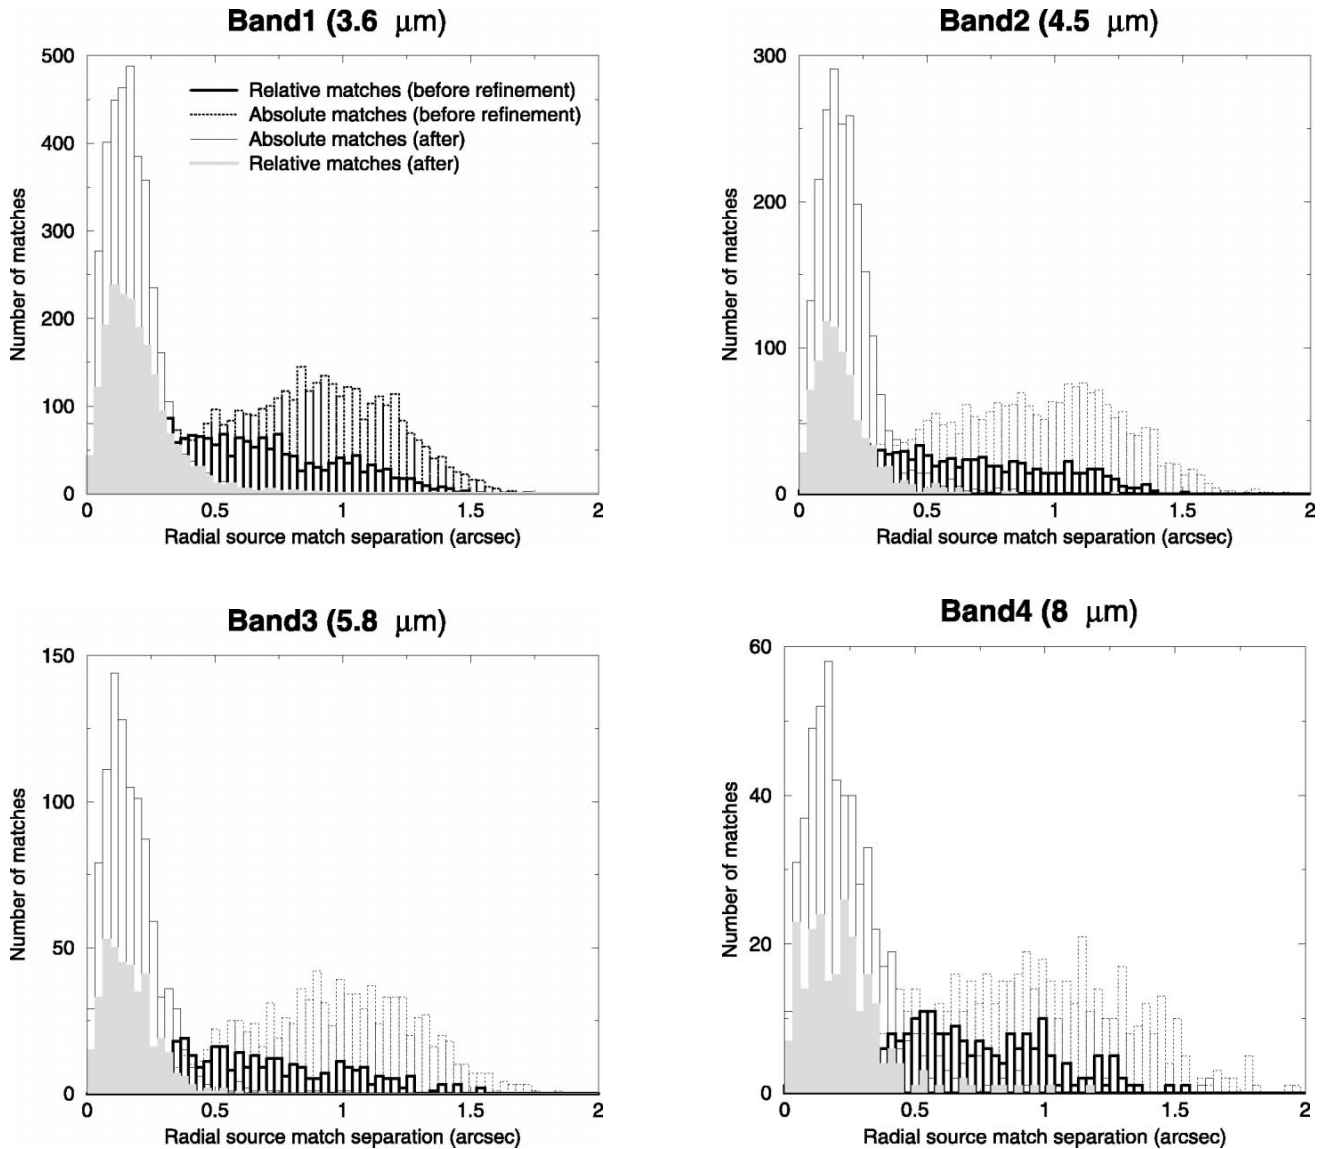


FIG. 9.—Distributions of matched source radial separations before and after refinement for all bands of *Spitzer*'s IRAC instrument. Relative and absolute astrometric (with 2MASS magnitudes $K_s \leq 15$) matches have been separated.

size of the PSF. The distribution in r at any given uncertainty cut σ_r can also be described by a Rayleigh distribution (see eqs. [16] and [17]). Furthermore, Figure 10 shows that typical systematic uncertainties due to inaccurately calibrated distortion are minimal and, if present, are expected to be much less than the centroiding errors. Guided by the simulation in § 3 and Figure 9, we conservatively conclude that the majority of image pointings must be refined to better ≈ 200 mas. In fact, we can predict the absolute dispersion in image pointings about truth using equation (13). Given the typical match statistics listed in Table 1, and centroiding errors of $0''.18$ and $0''.27$ for

bands 1 and 4, respectively, we expect dispersions of ~ 35 and ~ 140 mas (1σ radial) about truth for these bands, respectively. We expect these dispersions to be smaller once centroiding errors are brought down using better characterized PSFs.

In this case study, the number of astrometric (frame to absolute) matches are factors 1.5–2 times greater than relative (frame to frame) matches (see Table 1). This is because of the specific flux-difference thresholds used in source matching (4% and 50% for relative and absolute matches, respectively; see above). The results of Figure 9 assumed 2MASS point sources with $K_s \leq 15$. To ascertain the degree to which the number of

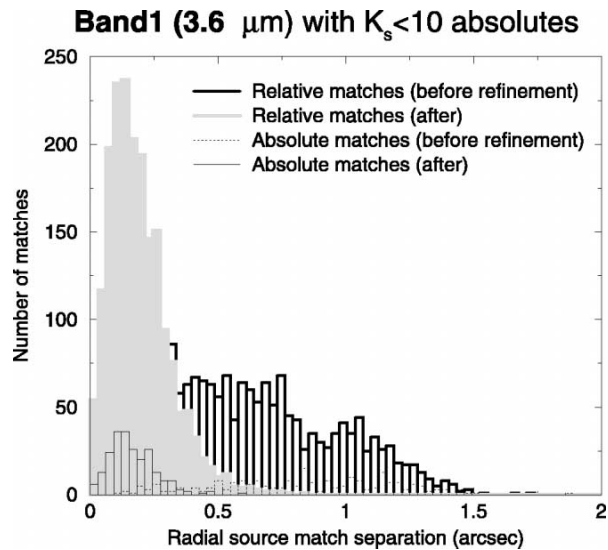



FIG. 10.—*Top*: Uncertainty in matched source radial separation (from centroid uncertainties of matches) as a function of actual source separation after refinement for IRAC band 1. *Bottom*: Same as Fig. 9 for IRAC band 1, but with astrometric matches having a 2MASS magnitude limit of $K_s = 10$.

absolute matches controls the level of refinement, we repeated the source matching and refinement for band 1 using a 2MASS magnitude cut of $K_s = 10$. The bottom panel of Figure 10 shows source separations before and after refinement. The number of absolutes per frame was randomly distributed between zero and four (compared to ~ 20 relative matches/frame). When frame-to-absolute separations after refinement are considered alone (*open histogram at bottom left*), the mean separation is smaller by a factor of ~ 2 , compared to the $K_s \leq 15$ case (Fig. 9, *top left panel*). This is a result of inherently lower positional

uncertainties for the brighter ($K_s \leq 10$) sources. When combined with relative matches however, the mean source separation and hence the level of refinement is essentially unchanged, compared to our results above.

5. DISCUSSION AND CONCLUSIONS

We have presented a generic algorithm to perform astronomical image registration and pointing refinement. It is generic in the sense that it can be used on any set of astronomical images that recognize the FITS and WCS pointing standards (Greisen & Calabretta 2002; Calabretta & Greisen 2002). Either relative (self-consistent frame-to-frame registration), single image absolute-astrometric, or simultaneous (relative and absolute) refinement is supported. The crux of the method involves matching point-source positions between overlapping image frames and using this information to compute image offset corrections by globally minimizing a weighted sum of matched point-source positional differences.

To ensure robust registration and refinement, the algorithm is best optimized with the following criteria:

1. The random uncertainty in measured twist angle of an individual image frame ($\delta\theta$) is assumed to be small so as to ensure $\sin \delta\theta \approx \delta\theta$. $\delta\theta \lesssim 60'$ is a good working measure for the intended applications of this algorithm [where $1 - (\sin \delta\theta)/\delta\theta \lesssim 10^{-4}$].

2. Input images have been accurately calibrated for distortion and possible nonuniform pixel scale. Any position-dependent systematic offset between source matches will limit the refinement accuracy to the size of the systematic error involved.

3. Sufficient area overlap between adjacent image frames is needed to ensure good match statistics.

4. Availability of *point* sources with well-defined flux distribution profiles approaching that of the instrument/detector's PSF. Extended sources will lead to larger centroiding errors.

5. Well-characterized PSF(s) for the image(s) at hand. These are crucial for accurate determination of source extraction centroids. If the inherent telescope pointing uncertainty is of an order of a third or larger than the detector pixel size, centroiding accuracies to better than one-tenth of a pixel or resolution element are recommended. For pointing uncertainties much less than the pixel size, there is little to be gained in resolution by improving the registration.

6. With the suggested centroiding accuracy from step 5, at least five relative (frame to frame) and five absolute source matches per frame will give pointings refined to better than three-hundredths of a pixel (rms). The greater the number of matches, the better the refinement. A minimum of two source matches per image frame, either relative, absolute, or both, is required to unambiguously determine all offset parameters per image.

7. Astrometric (absolute) reference sources should be used wherever possible. These provide a baseline to counteract any systematic deviations from "truth" or expected pointing in the ICRS, especially if absolute refinement is desired.

8. It is assumed that uncertainties in image pointing and point-source extraction centroids are random and independent. Prior image pointing uncertainties should be used, if available. These will prevent erroneous overcorrecting of the pointing for cases in which it is well known a priori.

Our simulations show that potentially good refinement can be obtained with minimal requirements. For a large fraction of images in an ensemble, refinements of better than ≈ 65 mas (1σ radial from truth) can be obtained, with an average of ~ 10 relative and ~ 20 absolute matches per frame with extraction centroids $\leq 0''.2$. This amounts to an improvement of 95% relative to truth for a majority of images. This could be better with higher match statistics and/or more accurate centroids, since the 1σ dispersion about truth typically scales as $\sim \sigma_{\text{ext}}(2/N_{\text{ext}})^{1/2}$ for a given centroiding uncertainty σ_{ext} and number of matches N_{ext} .

Analysis of observations from *Spitzer's* IRAC instrument shows that the dispersion in source separations after refinement is entirely consistent with the inherent dispersion in extraction centroid uncertainties. This implies that systematic uncertainties such as inaccurately calibrated distortions are negligible, since otherwise, dispersions in matched source separations after refinement would be larger relative to centroiding errors. Comparing dispersions of refined pointings about *truth* with matched source separations as found in our simulation, and rescaling to the appropriate numbers of matches using $1/N_{\text{match}}^{1/2}$ scaling, we predict (at the 2σ level) refinements to

better than ~ 70 and ~ 280 mas for IRAC bands 1 and 4, respectively. These bands bracket two extremes in available source matches, and these refinement estimates correspond to ~ 55 and ~ 8 (astrometric and relative) matches per band-dependent frame, respectively.

The goal of astronomical image registration is to exploit the resolution capabilities of existing and upcoming detectors whose pointing control and stability may not evolve at the same rate. The aim is to optimize the achievable S/N and science return therein. The algorithm presented here is just the tip of the iceberg for exploring one of many optimization techniques used in the diverse fields of image and signal processing and computer vision science.

F. J. M. is indebted to John Fowler for proofreading the manuscript and providing useful suggestions. We thank Howard McCallon for illuminating discussions, and David Shupe for assistance with simulations. This work is based in part on archival data obtained with the *Spitzer Space Telescope*, which is operated by the Jet Propulsion Laboratory, California Institute of Technology, under NASA contract 1407. Support for this work was provided by NASA through an award issued by JPL/Caltech. This publication makes use of data products from the Two Micron All Sky Survey, which is a joint project of the University of Massachusetts and the Infrared Processing and Analysis Center/California Institute of Technology, funded by NASA and the National Science Foundation.

APPENDIX A

ELEMENTS OF THE COEFFICIENT MATRIX

In this section, we provide general analytic expressions for elements of the coefficient matrix \mathbf{M} (eq. [10]). These elements are obtained by applying the minimization conditions (eq. [9]) to the cost function defined by equation (7). The “base” coefficient labels A , B , and C correspond to the three equations obtained by evaluating the partial derivatives in equation (9) (which we call A , B , and C from left to right, respectively). Position and uncertainty variables appearing in the expressions below were defined in § 2.2.

$$\begin{aligned}
 A_{\theta}^m &= \sum_n \sum_i \frac{(y_i^m - y_c^m)^2}{\Delta x_i^{m,n}} + \frac{(x_i^m - x_c^m)^2}{\Delta y_i^{m,n}} + \sum_n \frac{1}{\sigma_{\theta n}^2}, \\
 A_X^m &= - \sum_n \sum_i \frac{(y_i^m - y_c^m)}{\Delta x_i^{m,n}}, & A_Y^m &= \sum_n \sum_i \frac{(x_i^m - x_c^m)}{\Delta y_i^{m,n}} \\
 A_{\theta}^n &= - \sum_n \sum_i \frac{(y_i^m - y_c^m)(y_i^n - y_c^n)}{\Delta x_i^{m,n}} + \frac{(x_i^m - x_c^m)(x_i^n - x_c^n)}{\Delta y_i^{m,n}}, \\
 A_X^n &= \sum_i \frac{(y_i^m - y_c^m)}{\Delta x_i^{m,n}}, & A_Y^n &= - \sum_i \frac{(x_i^m - x_c^m)}{\Delta y_i^{m,n}}
 \end{aligned}$$

$$\begin{aligned}
 B_\theta^m &= A_X^m, & B_X^m &= \sum_n \sum_i \frac{1}{\Delta x_i^{m,n}} + \sum_n \frac{1}{\sigma_{X_m}^2}, & B_\theta^n &= \sum_i \frac{(y_i^m - y_c^m)}{\Delta x_i^{m,n}}, & B_X^n &= - \sum_i \frac{1}{\Delta x_i^{m,n}} \\
 C_\theta^m &= A_Y^m, & C_Y^m &= \sum_n \sum_i \frac{1}{\Delta y_i^{m,n}} + \sum_n \frac{1}{\sigma_{Y_m}^2}, & C_\theta^n &= - \sum_i \frac{(x_i^n - x_c^n)}{\Delta y_i^{m,n}}, & C_Y^n &= - \sum_i \frac{1}{\Delta y_i^{m,n}} \\
 \Psi_A^m &= - \sum_n \sum_i \frac{(y_i^m - y_c^m)(x_i^n - x_c^n)}{\Delta x_i^{m,n}} + \frac{(x_i^m - x_c^m)(y_i^m - y_c^m)}{\Delta y_i^{m,n}}, \\
 \Psi_B^m &= - \sum_n \sum_i \frac{(x_i^m - x_c^m)}{\Delta x_i^{m,n}}, & \Psi_C^m &= - \sum_n \sum_i \frac{(y_i^m - y_c^m)}{\Delta y_i^{m,n}}.
 \end{aligned}$$

APPENDIX B

THE ERROR-COVARIANCE MATRIX

The full error-covariance matrix is one of the byproducts of the POINTINGREFINE software. This reports all variances and covariances for and between all (inter and intra) image offsets that are necessary for refinement from the global minimization. It can be used to explore the strength of long-distance correlations between images in a mosaic and the presence of any undue systematic walks after refinement. The latter could arise if one lacks the desired number of absolute astrometric sources or the correct magnitude of prior image pointing uncertainties, as was discussed in § 2.3.

For a given pair of images (i, j) and three computed offsets per image ($\delta\theta, \delta X, \delta Y$), we have a total of six possible (correlated) offset pairs or covariance matrices. The possible covariances (or variances if $i = j$ for the same offsets) for any two images are therefore

$$\begin{aligned}
 &\text{Cov}(\theta_i, \theta_j), \text{Cov}(\theta_i, X_j), \text{Cov}(\theta_i, Y_j), \\
 &\text{Cov}(X_i, X_j), \text{Cov}(X_i, Y_j), \text{Cov}(Y_i, Y_j). \tag{B1}
 \end{aligned}$$

If we define any of these covariances generically as $\text{Cov}(\alpha_i, \beta_j)$ then the format for the error-covariance matrix for all possible image pair combinations (i, j), where ($i = 1, 2 \dots m; j = 1, 2 \dots m$), is

$$\begin{aligned}
 &\text{Cov}(\alpha_i, \beta_j) = \\
 &\begin{bmatrix} \text{Cov}(\alpha_1, \beta_1) & \text{Cov}(\alpha_1, \beta_2) & \dots & \text{Cov}(\alpha_1, \beta_m) \\ \text{Cov}(\alpha_2, \beta_1) & \text{Cov}(\alpha_2, \beta_2) & \dots & \text{Cov}(\alpha_2, \beta_m) \\ \dots & \dots & \dots & \dots \\ \text{Cov}(\alpha_m, \beta_1) & \text{Cov}(\alpha_m, \beta_2) & \dots & \text{Cov}(\alpha_m, \beta_m) \end{bmatrix}. \tag{B2}
 \end{aligned}$$

Figure 11 shows gray-scale representations of covariance matrices for all combinations of offset parameters (eq. [B1]) for the 105 image IRAC test case, assuming *both relative and*

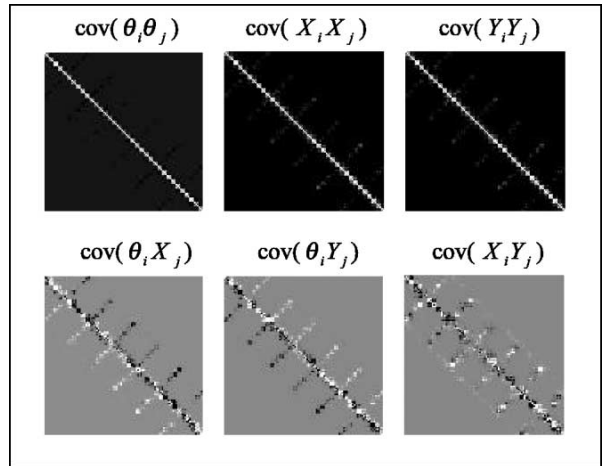


FIG. 11.—Gray-scale representation of covariance matrices for all orthogonal and rotational offsets in IRAC test case when absolute astrometric matches are used. See eq. (B2) for matrix definition.

absolute astrometric source matches in the refinement. With 105 images, each covariance matrix has 105×105 elements (or pixels, in this representation). It is important to note that the cross-correlation between different offset types is not symmetric. For instance, the correlation between θ_i and X_j is not the same as θ_j and X_i . Figure 12 shows the same set, but with *only relative* (frame to frame) matches used for the refinement. Comparing Figures 11 and 12, we draw the following conclusions: first, for the case including astrometric matches, “long distance” correlations in image offsets are greatly reduced. This is the result of astrometric sources anchoring each image to the fiducial reference frame, making them more independent of each other. The covariance matrices

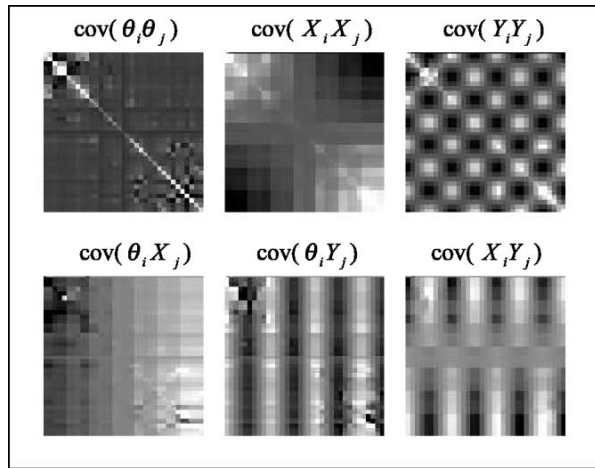


FIG. 12.—Gray-scale representation of covariance matrices for all orthogonal and rotational offsets in IRAC test case when absolute astrometric matches are *not* used (i.e., only relative frame-to-frame matches used). See eq. (B2) for matrix definition.

become essentially block-diagonal. Second, uncertainties (variances along the diagonal for the same offset pair combination) are greatly attenuated when astrometric matches are used compared to the relative-only match case. There are greater numbers of degrees of freedom per image when astrometrics are included, and this reduces the relative uncertainty.

The relative-only case exhibits greater long-distance correlations (larger off-diagonal values), since image positions are dictated solely by frame-to-frame matches in image overlap regions. This makes each successive image position dependent on its nearest neighbor positions, which depend on their own neighbors, and so on throughout the system of linked images.

Apart from astrometric matches reducing long-distance correlations, this can also happen if prior image pointing uncertainties are intrinsically smaller. Small priors will pull refinement offsets from the global minimization toward zero, which minimizes the $L_{\text{a priori}}$ term in equation (4). The priors force a constraint on each individual image to prevent “overrefinement,” regardless of the number of source matches present. Each image therefore becomes more independent of its neighbors, and long-distance correlations are reduced.

We also note the rich and diverse patterns in the covariance matrices for this observation, especially the relative-only case (Fig. 12). These are characteristic of the image layout and mosaic map geometry (see Fig. 8). Take, for instance, the $\text{Cov}(Y_i, Y_j)$ matrix in Figure 12. The checkerboard pattern arises from the relative image numbering in the map and how this translates to the numbering of elements in the covariance matrix (eq. [B2]). In the mosaic, the images repeat from right to left to create a leg, then left to right and down again, for a total of seven legs. Adjacent image pairs (i, j) along vertical sections in the map are strongly correlated in their (Y_i, Y_j) offsets, while widely separated images are less correlated. For instance, image 1 at top right has its Y offset strongly correlated with the Y offsets of images directly below it (i.e., image numbers 30, 31, 60, 61, 90, and 91). The high correlations are therefore with every ~ 30 th image giving bright regions in the gray-scale covariance image. In addition, the Y offset of image 1 is least correlated with that of images in the far left vertical strip (images 15, 16, 45, 46, 75, 76, and 105), resulting in dark regions in the covariance image.

In the end, one purpose of the covariance matrix is to visualize the degree of correlation between image positions in the ensemble as a whole. This allows one to ascertain whether refinement solutions are driven by any particular invalid input assumptions (e.g., priors), or insufficient astrometric reference source information, if robust absolute refinement is desired.

REFERENCES

- Barnea, D. I., & Silverman, H. F. 1972, *IEEE Trans. Comput.*, 21, 179
- Bly, M. J., et al. 2002, in *ASP Conf. Ser. 281, Astronomical Data Analysis Software and Systems XI*, ed. D. A. Bohlender, D. Durand, & T. H. Handley (San Francisco: ASP), 513
- Brown, L. G. 1992, *ACM Computing Surveys*, 24, 4, 325
- Bustos, I. H., & Calderón, J. H. 2000, *Rev. Mex. AA*, 39, 303
- Calabretta, M. R., & Greisen, E. W. 2002, *A&A*, 395, 1077
- Davis, L. E. 1996, in *ASP Conf. Ser. 101, Astronomical Data Analysis Software and Systems V*, ed. G. H. Jacoby & J. Barnes (San Francisco: ASP), 147
- Evans, M., Hastings, N., & Peacock, B. 2000, in *Statistical Distributions* (3rd ed.; New York: Wiley), 57
- Fazio, G. G. 2004, *ApJS*, 154, 258
- Greisen, E. W., & Calabretta, M. R. 2002, *A&A*, 395, 1061
- IPF Team. 2003, *SIRTF Instrument Pointing Frame Kalman Filter Execution Summary* (JPL Int. Rep. ID01M095; Pasadena: JPL)
- Kuglin, C. D., & Hines, D. C. 1975, *Proc. IEEE*, 163
- Lawrence, C. R., et al. 2000, *Operational Implications of the Time-dependent Pointing Behavior of SIRTF*, Ver. 1.00 (JPL Int. Rep. 2000 Aug. 22; Pasadena: JPL)
- Lonsdale, C. J., et al. 2003, *PASP*, 115, 897
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1999, *Numerical Recipes in C: the Art of Scientific Computing* (2nd ed.; Cambridge: Cambridge Univ. Press)
- Valdes, F. G. 1998, in *ASP Conf. Ser. 145, Astronomical Data Analysis Software and Systems VII*, ed. R. Albrecht, R. N. Hook, & H. A. Bushouse (San Francisco: ASP), 53
- Veran, J.-P., & Rigaut, F. J. 1998, *Proc. SPIE* 3353, 426
- Xu, C. K., Lonsdale, C. J., Shupe, D. L., Franceschini, A., Martin, C., & Schiminovich, D. 2003, *ApJ*, 587, 90
- Zacharias, N., McCallon, H. L., Kopan, E., & Cutri, R. M. 2003, in *IAU Joint Discussion 16, The International Celestial Reference System: Maintenance and Future Realization* (Paris: IAU), 43