HIGH RESOLUTION IMAGE CONSTRUCTION FROM IRAS SURVEY – PARALLELIZATION AND ARTIFACT SUPPRESSION

Yu Cao and Thomas A. Prince
Division of Physics, Mathematics and Astronomy,
California Institute of Technology,
Pasadena, CA91125, USA

ABSTRACT.

The *Infrared Astronomical Satellite* carried out a nearly complete survey of the infrared sky, and the survey data are important for the study of many astrophysical phenomena. However, many data sets at other wavelengths have higher resolutions than that of the co-added *IRAS* maps, and high resolution *IRAS* images are strongly desired both for their own information content and their usefulness in correlation studies.

The HIRES program was developed by the Infrared Processing and Analysis Center (IPAC) to produce high resolution ($\tilde{}$ $1'$) images from *IRAS* data using the Maximum Correlation Method (MCM). In this paper, we describe the port of HIRES to the Intel Paragon, a massively parallel supercomputer. A speed increase of about 7 times is achieved with 16 processors and 5 times with 8 processors for a $1^\circ \times 1^\circ$ field. Equivalently a 64 square degree field can be processed using 512 nodes, with a speedup factor of 320.

Images produced from the MCM algorithm sometimes suffer from visible striping and ringing artifacts. Correcting detector gain offsets and using a Burg entropy metric in the restoration scheme were found to be effective in suppressing these artifacts.

1. Introduction

The *Infrared Astronomical Satellite* (*IRAS*) provided our first comprehensive look at the infrared sky, producing a nearly complete survey at mid- to far-infrared wavelengths (12, 25, 60, and 100 microns). Information about *IRAS* relevant to this paper is given in Section 3..

The Maximum Correlation Method (MCM) algorithm [1] produces high resolution images from the survey and additional observation (AO) data, using a nonlinear iterative scheme. The resulting images have resolution of about $1'$, compared to the $4' - 5'$ subtended by the 100 μm band detectors in the *IRAS* focal plane. Application of the algorithm to the *IRAS* data has been limited largely by the computational resources available for HIRES processing. A description of the MCM algorithm is outlined in Section 4..

We have ported the HIRES program to the Intel Delta and Paragon systems. Each $1^\circ \times 1^\circ$ field is mapped to an 8- or 16-node process grid, which shares the computation by loading different observation scans. An efficiency of 60 % is reached with 8 nodes. Section 2. further explains the motivation for this port, and Section 5. discusses the parallelization strategy, output verification, and performance analysis.

In Sections 6. and 7. we offer descriptions of artifact reduction algorithms, namely using estimates of gain offset to eliminate striping, and using a Burg entropy prior in the iterative algorithm to suppress ringing around bright point sources.

2. Scientific Motivation

The wavelength bands covered by the *IRAS* survey are tracers of star-forming regions and numerous other components of the interstellar medium. A variety of studies have been made to date ranging from structure on a galactic scale to detailed studies of individual molecular clouds (see [2, 11]). The strength of *IRAS* is the completeness of the survey. However, in many cases the spatial resolution of the comparison data sets at other wavelengths is better than for *IRAS*, and thus the $4' - 5'$ resolution of the released *IRAS* images (the Infrared Sky Survey Atlas, ISSA) sometimes limits the comparison. The desire for higher spatial resolution combined with the paucity of new infrared satellite missions has inspired many efforts to extract high spatial resolution information from the data (e.g. [3, 4]). The products most widely accessible to the science community are the HIRES images distributed by the Infrared Processing and Analysis Center (IPAC), which are based on the Maximum Correlation Method.

The HIRES process is very demanding computationally. A $1^\circ \times 1^\circ$ field of typical scan coverage takes 1 – 2 hours of CPU time on a Sun SPARCstation 2, for all four wavelength bands and 20 iterations (at which point artifacts limit further improvement of image quality).

As part of a program in high-performance computational science and engineering, Caltech has developed significant software and hardware capabilities for massively parallel computing. The high demand for HIRES images, along with the availability of parallel computing facilities, motivated the port of HIRES to the parallel supercomputers.

We also developed algorithms which can effectively suppress the artifacts, which allows the iteration procedure to be carried much further (hence requiring more CPU time and further justifying the parallel computing approach).

3. Relevant Information about *IRAS*

The *IRAS* focal plane was designed for the identification of point sources. It included eight staggered linear arrays subtending $30'$ in length, two in each of four spectral bands at 12, 25, 60, and 100 μm. Data rate considerations forced the detector sizes to be much larger than the diffraction limit of the telescope. The typical detector sizes were 45×267, 45×279, 90×285, and 180×303 arcsec (full width at half maximum response, FWHM) respectively, at the four wavelength bands. The sky was scanned in "push-broom" fashion.

This combination of focal place, detector size, and scan pattern optimized detection of point sources in areas of the sky where the separation between sources was large compared to the sizes of the detectors. However, it complicated the construction of images of regions containing spatial structure on the scale of arcminutes.

4. The Maximum Correlation Method

Starting from a model of the sky flux distribution, the HIRES MCM algorithm folds the model through the *IRAS* detector responses, compares the result track-by-track to the observed flux, and calculates corrections to the model. The process is taken through about 20 iterations at which point artifacts limit further improvement. The algorithm yields a resolution of approximately $1'$ at 60 μm. This represents an improvement in resolution by as much as a factor of 20 in solid angle over the previous images from the *IRAS* Full

Resolution Survey Coadder (FRESCO). We give a brief description of the MCM algorithm following the formalism and notations of [1].

Given an image grid f_j, with n pixels $j = 1, ..., n$ and m detector samples (footprints) with fluxes

$$D_i : i = 1, ..., m, \tag{1}$$

whose centers are contained in the image grid, an image can be constructed iteratively from a zeroth estimate of the image, $f_j^0 = \text{const.} > 0$ for all j. In other words the initial guess is a uniform, flat, and positive definite map. For each footprint, a correction factor C_i is computed as,

$$C_i = D_i / F_i, \tag{2}$$

where

$$F_i = \sum_j r_{ij} f_j, \tag{3}$$

and r_{ij} is the value of the ith footprint's response function at image pixels f_j. Therefore F_i is the current estimate of the ith footprint's flux, given image grid f_j.

A mean correction factor for the jth image pixel is computed by projecting the correction factor for the footprints into the image domain:

$$c_j = \left[\sum_i (r_{ij}/\sigma_i^2) C_i\right] / \left[\sum_i (r_{ij}/\sigma_i^2)\right]. \tag{4}$$

The weight attached to the ith correction factor for the jth pixel is r_{ij}/σ_i^2, where σ_i is the *a priori* noise assigned to the ith footprint.

The kth estimate of the image is computed by

$$f_j^{(k)} = f_j^{(k-1)} c_j. \tag{5}$$

In practice when the footprint noise σ_i is not easily estimated, an equal noise value for all footprints is assumed, and the MCM is identical to the Richardson-Lucy algorithm [10, 8].

5. Parallelization

Detector data are stored in scanlines called *legs*, which contain individual *footprints*. Profiling a typical HIRES process showed that more than 95 % of the total execution time was spent within the code which calculates the footprint and image correction factors. In the parallel decomposition of the problem, each processor takes care of footprints from a set of scanlines. The reasons for doing this are:

1. Small programming effort. The essence of the original HIRES architecture is left untouched.

2. Footprints in one leg share the same response function grid, except for a translation, which is basically the reason the original code processes the data one leg at a time. Keeping

Table 1: Speed comparisons for 60 μm band of M51

Sun SPARCstation 2	720 sec
Single node of the Paragon	640 sec
8 nodes of the Paragon	137 sec

the whole leg in one processor is therefore a natural choice, which minimizes local memory usage.

3. As we will discuss in Section 6., *IRAS* detectors have gain differences which are especially prominent for the 60 and 100 μm bands. The gain offset can be estimated from correction factors in the same leg, which came from the same detector.

Each node calculates the correction factor C_i's for its share of footprints, and projects them onto the pixels covered by the footprints. A global sum over all processors for the correction factor c_j's for each image pixel is performed at end of each iteration, and the weighted average is taken, which is then applied to the image pixel value.

Decomposition in the image pixel domain was not carried out for the $1^\circ \times 1^\circ$ field, eliminating the need for ghost boundary communication, which would be significant and complicated to code, due to the large size and irregular shape of the detector response function. This helped maintaining the parallel code similar in structure to the sequential one, making simultaneous upgrades relatively easy.

The efficiency of the parallel program depends on the scan coverage of the field processed. The computation time is roughly proportional to the total coverage (i.e. total number of footprints), while the communication overhead is not related to footprints and is only dependent upon the image array size. So the efficiency is higher for a field with higher coverage.

For a large field (e.g. $6^\circ \times 6^\circ$ of ρ Ophiuchus), the detector measurements are broken into $1^\circ \times 1^\circ$ pieces with overlap 0.15°. Each $1.15^\circ \times 1.15^\circ$ field was loaded on to a subgroup of 8 or 16 processors. The overlap was chosen to be large enough so that cropping the overlap after HIRES ensures smoothness at the boundaries. Therefore no inter-subgroup communication was needed during HIRES, at the cost of a moderate increase in computation.

The output images from the parallel computers are compared with those from the standard HIRES program running on a Sun SPARCstation. The differences are well within the range of numerical round-off errors. At the 20th iteration, the standard deviation of (NewImage - OldImage) / OldImage averages to about 10^{-4}.

The executable code was compiled and linked with a math library conformant to the IEEE 754 standard. For the 60 μm band of M51 (baseline removed data), a time comparison is shown in Table 1.

Efficiency is 60 % for 8 nodes for a $1^\circ \times 1^\circ$field. All 512 nodes can be used to process a 64 square degree field with a speedup factor of 320. The global sum operation, which collects pixel correction factors from different nodes, is the primary source of overhead in the parallel program.

6. Destriping Algorithm

Stripes are the most prominent artifacts of the HIRES images. HIRES takes in the *IRAS* detector data, and if not perfectly calibrated, would try to fit the gain differences in the detectors by a striped image. The striping builds up in amplitude and sharpness along with the HIRES iterations, as the algorithm refines the "resolution" of the stripes (see Fig. 1(a) and (b)).

The IPAC program LAUNDR [5] invokes several one dimensional flat fielding and deglitching techniques. For the purpose of destriping, the one dimensional algorithm works well for regions with a well-defined baseline, but the result is not satisfactory for regions where structure exists in all spatial frequencies.

Another IPAC utility KESTER, developed by Do Kester, is similar in principle to the approach we take. The idea is to process the data with HIRES to certain iterations to obtain an image, which is then used to simulate a set of detector flux measurements. The baseline offsets of the original data are then calibrated using the simulated data set.

Our approach is to combine the image construction and the destriping process. Since the striping gets amplified through the iterations, the idea of applying constraints to the correction factors is natural.

Assume footprints in the same leg L suffer from the same unknown gain offset G_L, then

$$D_i^* = G_L D_i \tag{6}$$

is the "true" detector flux, had the detector gain been perfectly calibrated. The G_L's can be seen as extra parameters to be estimated, besides the image pixels f_j. Under a Poisson framework, the maximum likelihood estimate for G_L is

$$\prod_{i \text{ in leg } L} \left(\frac{G_L D_i}{F_i} \right)^{D_i} = \prod_{i \text{ in leg } L} (C_i^*)^{D_i} = 1. \tag{7}$$

in which C_i^* is the gain compensated correction factor.

This choice of the unknown G_L minimizes the mutual information between the sets D_i^* and F_i in the leg, i.e. the resulting correction factors C_i^* will extract the minimum amount of information from the stream D_i^*. According to the maximum entropy principle, this is the only reasonable choice.

From another point of view, this strategy works because the procedure of averaging C_i's to get c_j has a smoothing effect on the image, so that the image f_j therefore F_i does not contain as much striping power as the footprints D_i, especially on the scale smaller than one detector size.

When the legs do contain non-random gain deviation roughly periodic on a scale larger than the detector size (typically around $7'$ for *IRAS*, distance between neighboring detectors), this destriping method sometimes fails to smooth out the wide stripes. It does eliminate high spatial frequency stripes, but may result in wide regions with discontinuous flux levels. A heuristic analogy for understanding this behavior can be made with the one dimensional Ising model where the energy is lowest when all the spin directions are lined up, if we compare the negative likelihood to the energy in the Ising model, and the residual gain offset to the spin. Just like the spins can be trapped in a local minimum energy state (aligned in patches of up's and down's), the gain estimation may reach a local maximum of

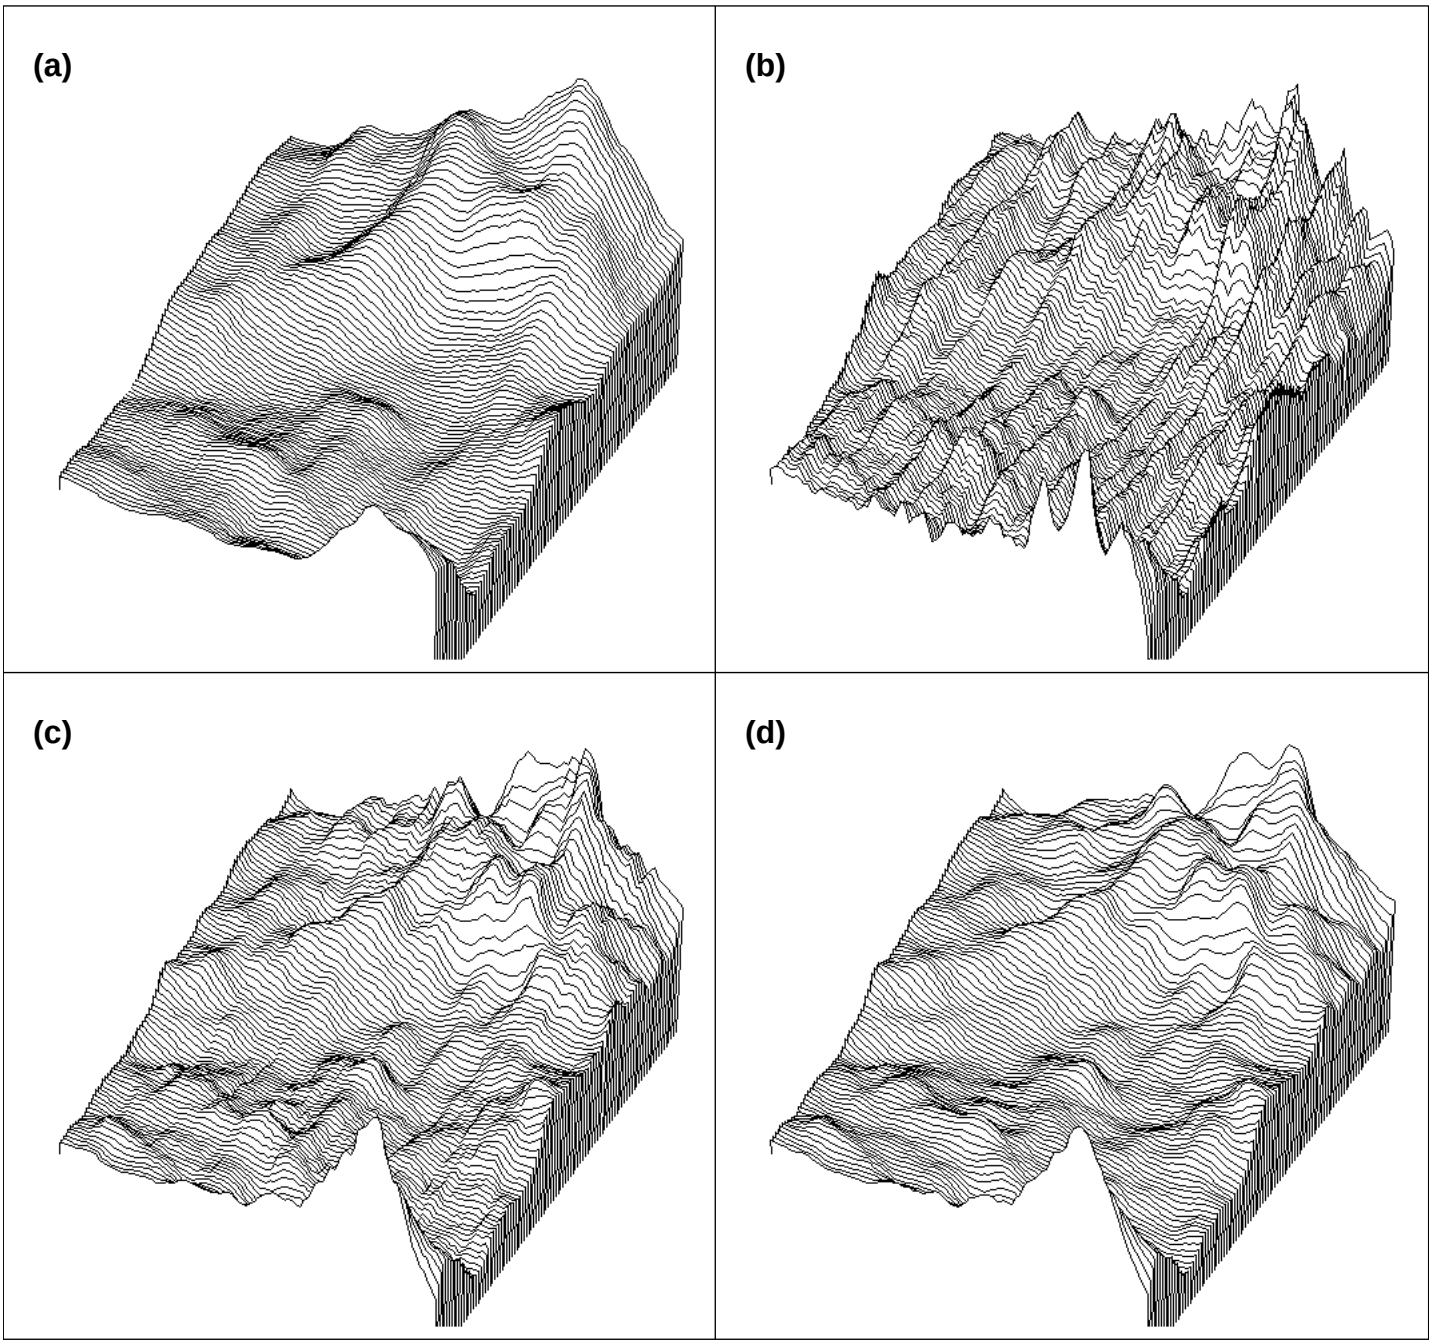

Figure 1: (a). 1st iteration image for a field in ρ Ophiuchus (100 μm band); (b). 20th iteration, standard HIRES; (c). 20th iteration, with uniform gain compensation; (d). 20th iteration, with local gain compensation. Size of image is $1^\circ \times 1^\circ$. Height of surface represents flux.

the likelihood function in the G_L space (image would consist of smooth regions at different flux levels). The LAUNDR program, which is run upstream of the MCM process, is capable of detecting this kind of flux variation and correcting it. But when raw data are fed into the MCM program without the LAUNDR offsets, it is necessary to first smooth the image f_j with a large kernel ($15'$), before trying to estimate the gain offsets.

A further complication lies in the fact that the assumption of a uniform gain offset in a certain leg is only approximately true. Various hysteresis effects cause the gain to drift slightly within the 1° range. The more aggressive form of the destriping algorithm estimates the gain offset locally as the geometric mean of the correction factors for nearby footprints, so the estimated gain correction for each footprint varies slowly along the leg. The local gain offset is not allowed to differ by more than 10 % from the global value, since the gain is not expected to drift that much over a 1° scale, and the variation in computed offset average is most likely due to real local structure. We used an averaging length of $10'$ to estimate the local offset. Because it is larger than the spatial resolution of the first iteration image ($5'$), it is safe to refer the average correction factor on that scale as due to gain offset. The $10'$ length scale is also small enough to capture the drifting behavior of the gain, as shown by visual inspection of output images as well as Fourier power spectrum analysis. Unlike the standard HIRES algorithm (in which stripes are amplified throughout the iterations), the local gain compensation decreases the striping power monotonically to a negligible level after roughly 10 iterations.

Fig. 1 demonstrates the striking effect of the destriping algorithm. Fig. 1(a) shows the first iteration image for a $1^\circ \times 1^\circ$ field in ρ Ophiuchus, which is smooth (blurry). Fig. 1(b) is the 20th iteration image of the field obtained with the standard HIRES algorithm, and is contaminated with strong striping artifacts. A tremendous improvement is seen in Fig. 1(c) which is produced with uniform gain compensation, although some weak stripes are still visible. Finally, using the local gain compensation method gives a stripe-free image, Fig. 1(d). It is also apparent that Fig. 1(d) contains many high spatial frequency features that are absent in 1(a).

7. De-ringing Algorithm

For many deconvolution problems, ringing artifact (or "ripple") appears when a bright point source exists over a non-zero background. The mechanism of the artifact can be understood as the Gibbs phenomenon (a sharp cutoff in high spatial frequency signal incurs ripples in the position domain).

A variant of the Log-Entropy MART [9]

$$f_j^{(k)} = f_j^{(k-1)} + (f_j^{(k-1)})^2 \sum_i \frac{r_{ij}}{F_i^2}(D_i - F_i) \tag{8}$$

was tested on *IRAS* data and gave satisfactory suppression of ringing artifact (Fig. 2).

The $(f_j^{(k-1)})^2$ factor in the correction term indicates a Burg entropy metric in the image space, and effectively boosts the correction factor for brighter pixels. So the bright point source is fitted better in the earlier iterations, which circumvents the corruption of background caused by the misfit.

The prior knowledge signified by using maximum Burg entropy estimation rule has

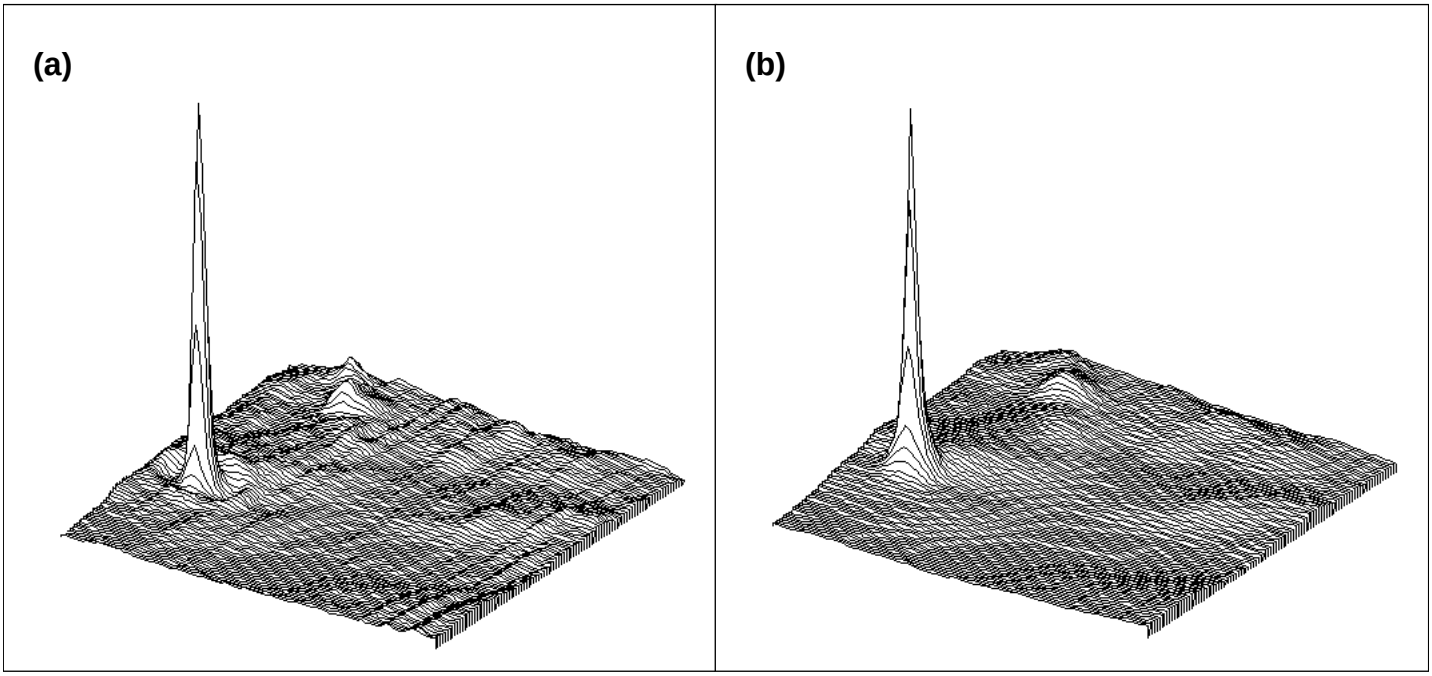

Figure 2: (a). Point source I16293-2422 in ρ Ophiuchus (100 micron), no ringing suppression; (b). Same field, using entropy prior for ringing suppression. Size of image is $1^\circ \times 1^\circ$. Peak flux in (a) is 3749 MJy/ster, and 3329 MJy/ster in (b).

been discussed in [7, 6]. According to [6], the class of optical objects described by the Burg entropy prior would tend to consist of a relatively small number of randomly placed bright cells, the rest being dim, befitting the bright point source scene we're concerned with.

Suppression of ringing leads to better photometry determination of the point source, and helps solve the source confusion problem, which is especially prominent in the Galactic plane.

8. Summary

The parallelization and algorithmic enhancements of the IPAC HIRES program have been described. These efforts will soon enable production of HIRES images by IPAC from the Intel Paragon supercomputer.

It is now possible to produce complete maps of the Galactic plane ($\pm 5^\circ$ latitude) at 60 and 100 μm with arcminute resolution, as well as maps of the Orion, Ophiuchus, and Taurus-Auriga clouds complexes. These maps will represent a 20-fold improvement in areal information content over current *IRAS* 60 and 100 μm maps and will be valuable for a wide range of scientific studies, including:

- The structure and dynamics of the interstellar medium (ISM)
- Cloud core surveys within giant molecular clouds
- Determination of initial mass functions (IMFs) of massive stars
- Study of supernova remnants (SNRs)

Additional information will come from combining the 60 and 100 μm HIRES data with the images and catalogs being produced from the 12 and 25 μm *IRAS* data by the Air Force Phillips Laboratory and Mission Research Corporation.

This research was performed in part using the Intel Touchstone Delta and the Intel Paragon operated by Caltech on behalf of the Concurrent Supercomputing Consortium.

We thank Tom Soifer, Joe Mazzarela, Jason Surace, Sue Terebey, John Fowler, Michael Melnyk, Chas Beichmann, Diane Engler and Ron Beck for their contributions to this project. YC also thanks Professor John Skilling for discussions during the workshop.

References

[1] H. H. Aumann, J. W. Fowler and M. Melnyk, "A Maximum Correlation Method for Image Construction of *IRAS* Survey Data," *Astronomical Journal*, Vol. 99(5), pp: 1674–1681, 1990.

[2] C. A. Beichman, "The *IRAS* View of the Galaxy and the Solar-System," *Annual Review of Astronomy and Astrophysics*, Vol. 25, pp: 521–563, 1987.

[3] T. R. Bontekoe, D. J. M. Kester, S. D. Price, A. R. W. Dejonge and P. R. Wesselius, "Image Construction from the *IRAS* Survey," *Astronomical Journal*, Vol. 248(1), pp: 328–336, 1991.

[4] T. R. Bontekoe, E. Koper and D. J. M. Kester, "Pyramid Maximum-Entropy Images of *IRAS* Survey Data," *Astronomy and Astrophysics*, Vol. 284(3), pp: 1037–1053, 1994.

[5] J. W. Fowler and M. Melnyk, *LAUNDR Software Design Specifications*, IPAC, Caltech, 1990.

[6] B. R. Frieden, "Estimating Occurrence Laws with Maximum Probability, and the Transition to Entropic Estimators," in *Maximum-Entropy and Bayesian Methods in Inverse Problems*, eds. C. R. Smith and W. T. Grandy, Jr., pp: 133–170, D. Reidel Publishing Company, Dordrecht, Holland, 1985.

[7] E. T. Jaynes, "Monkeys, Kangaroos, and N," in *Maximum Entropy and Bayesian Methods in Applied Statistics*, ed. J. H. Justice, pp: 26–58, Cambridge University Press, 1986.

[8] L. B. Lucy, "An Iterative Technique for the Rectification of Observed Distributions," *Astronomical Journal*, Vol. 79, pp: 745–754, 1974.

[9] A. R. De Pierro, "Multiplicative Iterative Methods in Computed Tomography," in *Mathematical Methods in Tomography*, eds. G. T. Herman, A. K. Louis and F. Natterer, pp: 167–186, Springer-Verlag, 1991.

[10] W. H. Richardson, "Bayesian-Based Iterative Method of Image Restoration," *Journal of the Optical Society of America*, Vol. 62, pp: 55-59, 1972.

[11] B. T. Soifer, J. R. Houck and G. Neugebauer, "The *IRAS* View of the Extragalactic Sky," *Annual Review of Astronomy and Astrophysics*, Vol. 25, pp: 187–230, 1987.