Infrared nonredundant mask imaging at Palomar

W. N. Weir, A. M. Ghez, P. W. Gorham, C. A. Haniff,
S. R. Kulkarni, K. Matthews, G. Neugebauer

California Institute of Technology, Division of Physics, Mathematics and Astronomy
Pasadena, California, 91125

ABSTRACT

We have recently begun a program of high resolution infrared imaging at Palomar Observatory. As a part of this effort, we have investigated the use of nonredundant masks both as an imaging technique and as a method of analyzing various aspects of our imaging system. In particular, we have applied the technique to a bright star and binary system using a three-hole mask. We find the method is useful for understanding certain systematic biases in our data, as well as in producing high quality images despite sparse uv coverage. We find that the use of multi-r_o apertures along with a large bandwidth does not significantly hamper image reconstruction, but provides significant extra coverage in the uv plane.

1. INTRODUCTION

The objectives of our infrared nonredundant mask (NRM) imaging program at Palomar are manifold. One goal is to understand the systematic effects of our imaging system. In contrast to fully filled aperture (FFA) imaging, NRM makes such effects more tractable by limiting the uv coverage and the portion of the mirror in use for any given observation. We are also interested in optimizing the NRM technique as an imaging tool in this new wavelength regime. At optical wavelengths, NRM has been shown to be superior to FFA as an imaging method for bright sources[1,2]. Haniff et al.[3] and Nakajima et al.[2] previously demonstrated the successful use of this technique in high resolution optical imaging. Infrared imaging significantly differs from the optical case in that the Fried parameter, r_o, characterizing the coherent patch size, is larger and, for existing detectors, read-out noise dominates over the photon noise. One must understand and explicitly allow for these factors to make most efficient use of masks. Possibly more importantly, however, NRM imaging is a small scale analogue of the use of long baseline discrete interferometers, so a systematic investigation of this technique may contribute to a better understanding of larger scale systems to be completed in the near future.

2. OBSERVATIONS AND PRELIMINARY REDUCTION

The data described herein were obtained during observing runs in August 1989, on the Hale 200 inch telescope at Palomar Observatory. Our detector is a Santa Barbara Research Center InSb array (58 by 62 pixels) placed at the f/415 Gregorian focus. The scale of the array is 55 ± 1 mas/pixel. All of our images were taken in the K band (2.2 μm, $\frac{\Delta\lambda}{\lambda} \sim 20\%$), and are the result of using a three-hole mask placed just below the secondary mirror. Each aperture in the mask is located at the corner of an equilateral triangle centered upon the optical axis of the telescope. The distance between the centers of each aperture in the mask is 33.5 mm, corresponding to 2.85 m on the primary mirror. The diameter of an aperture is 6.0 mm, or 0.51 m on the primary. The value of r_o at the time of these observations, as estimated from speckle data obtained on the same night, was approximately 0.3 m. Thus, despite our original intentions, these observations were actually multi-r_o in nature.

For each object or sky field, we took sets of 400 image frames of 100 milliseconds each. Our procedure was to integrate on the object of interest, the calibrator star, and the sky before changing the position angle of the mask and repeating the three sets of exposures.

The observations we have chosen to analyze in this paper are those made of SAO 86224, a binary star system of separation ~ 660 mas, and the calibrator point source SAO 86142. The estimated m_K of the binary is approximately 2.3, while the calibrator has $m_K \sim 3.6$. We purposefully restrict ourselves to examining the data from these objects first because of their simple morphology, allowing us to compare relatively accurate models with the data obtained.

We follow standard methods in the initial reduction of our image stacks. We first subtract individually scaled median filtered sky images obtained from the stack of 400 sky frames taken nearest in time to the object or calibrator frames

being reduced. Next we flatfield the data by dividing it by the difference between a median filtered sky flat and a dark frame. We remove dead and hot pixels in the image by interpolating over them. Finally, we pad the images out to 64 by 64 for computing fast Fourier transforms.

Fourier amplitudes and phases are recovered by separate processes. The amplitudes are obtained by computing power spectra for each of the frames from all stacks: object, calibrator and sky. The average power spectrum of the sky, representing the noise power spectrum, is subtracted from the power spectrum of each calibrator and object frame. The average of these noise subtracted power spectra are then computed. The square root of the normalized division of these two power spectra is then taken to be the sampled visibilities of the object. We describe the results of this stage of reduction in the next section. Phase recovery is accomplished through bispectral, equivalently closure phase, analysis. We investigate this process in section four.

3. AMPLITUDE ANALYSIS

Figure 1a is a grayscale plot of a single, typical 100 millisecond frame of the calibrator source using the three-hole mask. The signal is relatively strong but the overlapping fringes are barely distinguishable. Figure 1b is the co-added power spectrum of a stack of 400 such frames. Such a plot reveals many important features of our data.
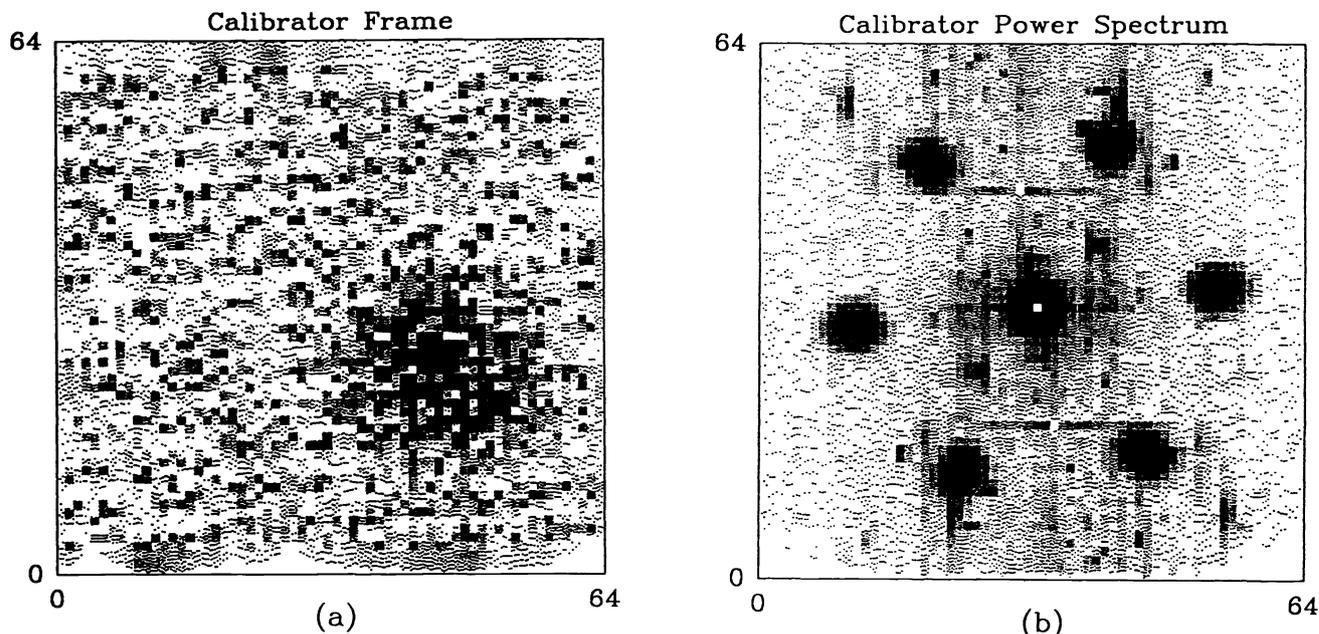

Figure 1. A typical 100 millisecond image of the calibrator source using the three-hole mask (a) and the co-added power spectrum of a stack of 400 such frames (b).

First note that there is a relatively uniform bias in the power spectrum. Simple analysis reveals that the estimated bias in the power spectrum, P_{bias}, due to a uniform readout noise across the chip, σ_r, should be

$$P_{bias} = \sigma_r^2 p$$

where p is the number of pixels in the image. The value of σ_r estimated for our array by analyzing several stacks of image frames is approximately 440 electrons, giving an estimated value of $P_{bias} \simeq 7.0 \times 10^8$. A random sampling of the power in several of the uncovered uv points in the calibrator power spectrum yields a mean value of 6.5×10^8, in very good agreement with the prediction. Subtracting the noise power spectrum from the data removes this bias almost entirely.

There are several points and areas in the uv plane that have significant power due to noise. Some of these noise features we have come to understand as artifacts of the detector readout electronics. Most of these anomalies in the power spectrum remain fairly constant over time, subtracting well from the data. Other features seem to be a bit too variable, however, and we simply choose to avoid using those areas of the uv plane when reconstructing images.

We are also interested in how the signal to noise ratio (SNR) of the power spectrum behaves as a function of the number of frames collected. In particular, we wish to know whether reading out the array at very high rates introduces correlations or anticorrelations in the data. Figure 2 is a plot of of the SNR of a peak point in the power spectrum of the calibrator versus the number of images used. Note that the SNR increases as the square root of the number of images up to 400 images.
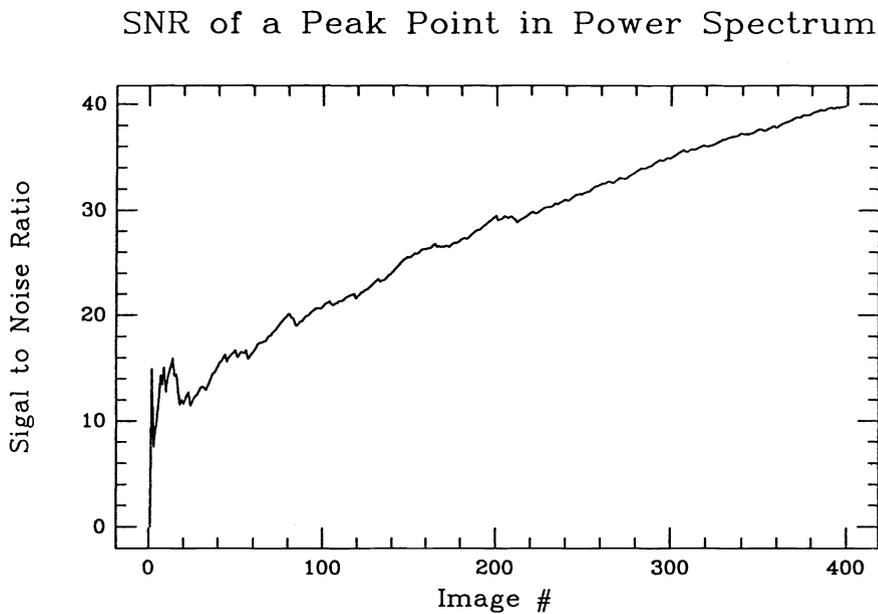
SNR of a Peak Point in Power Spectrum

Figure 2. The SNR of a peak point in the power spectrum of the calibrator source as a function of the number of images used. Note that the SNR increases at the square root of the number of images.

Finally, we wish to call attention to the the broad distribution of power about the peaks in the uv plane. This breadth is due to both the finite size of the apertures in the mask as well as the large bandwidth used. An important question is what fraction of these areas in the uv plane are actually useable for amplitude and phase recovery. Figure 3 demonstrates that at least for deriving visibilities, very large portions of the uv plane are reliably sampled for as few as two rotations of the mask. Figure 3 represents the sum of the visibilities obtained for two rotations of the mask for the binary star SAO 86224. In this case, the uv region covered by just a single baseline is large enough to contain a complete cycle of the sinusoid pattern expected in the power spectrum of a binary. In fact, the image parameters of the binary system can be almost completely determined using the visibilities from just these two rotations of the three-hole mask.

4. PHASE ANALYSIS

We employ bispectral analysis to recover the phases in our images. Since for this set of data we are principally interested in learning more about the systematics of our imaging system, our primary concern in analyzing the bispectrum has been to understand its SNR properties and any biases. A particular aim of our NRM program is to better understand the systematic biases found in the bispectrum of our FFA data. By using a mask, it is much easier to physically isolate the effects that may be biasing a particular closure phase.
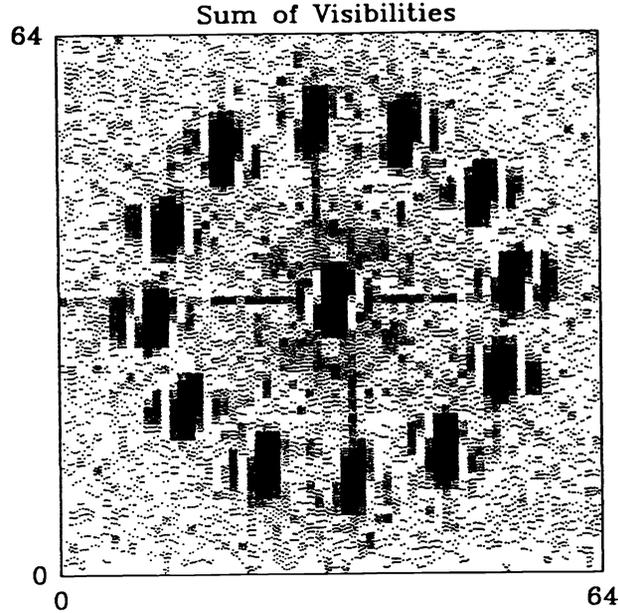
Sum of Visibilities

Figure 3. The sum of the visibilities for the binary star SAO 86224 for two rotations, 30° apart, of the three-hole mask. Note the relatively large fraction of the uv plane covered by just a single baseline.

Figure 4 is an Argand diagram of individual bispectrum phasors for a stack of calibrator frames. In this example, the bispectrum point being analyzed corresponds to the equilateral triangle connecting the points at the very center of each aperture. Because such a triangle is the most redundant for the given geometry, each of the corresponding uv points are local maxima within the power spectrum and, as one might expect, they define the triangle with the highest bispectrum SNR.

It is important here to note our definition of the SNR of a bispectrum phasor, as such a quantity is not uniquely defined for a complex number. Since we are only interested in deriving phases from the bispectrum, it is natural that we choose a definition that only depends upon the variance of the phase of a bispectrum phasor. Such a measure is especially important since the errors in the bispectrum differ quite significantly from a normal distribution. As we see in Figure 4, the variance in the modulus of the bispectrum phasor can be quite high while the variance in the phase remains low. The definition we employ is that suggested by Ayers[4], whereby one measures the variance in the direction perpendicular to the mean bispectrum phasor, B. This quantity, σ_p^2, can be computed using values of the variance in the real and imaginary parts of the phasor (σ_R^2 and σ_I^2, respectively), and the angle Φ measured counter-clockwise from the real axis to the mean bispectrum phasor:

$$\sigma_p^2 = \sigma_R^2 \sin^2 \Phi + \sigma_I^2 \cos^2 \Phi - cov(\text{Re}(B)\text{Im}(B)) \sin 2\Phi.$$

The SNR is then

$$SNR = \frac{|<B>|}{\sigma_p} \cdot \sqrt{n}$$

where n is the number of images used. Using this measure, we compute the SNR of the mean bispectrum phasor for three different points within the bispectrum space as a function of the number of images used (Figure 5). Note that just as in the case of the power spectrum, the SNR of the bispectrum phasors seems to increase without bound as the square root of the number of images.

Figure 6 is a plot analogous to figure 4, but involving a different effective triangle of baselines between the apertures. The uv points corresponding to these baselines lie just off the peaks in the power spectrum. The result is an ensemble
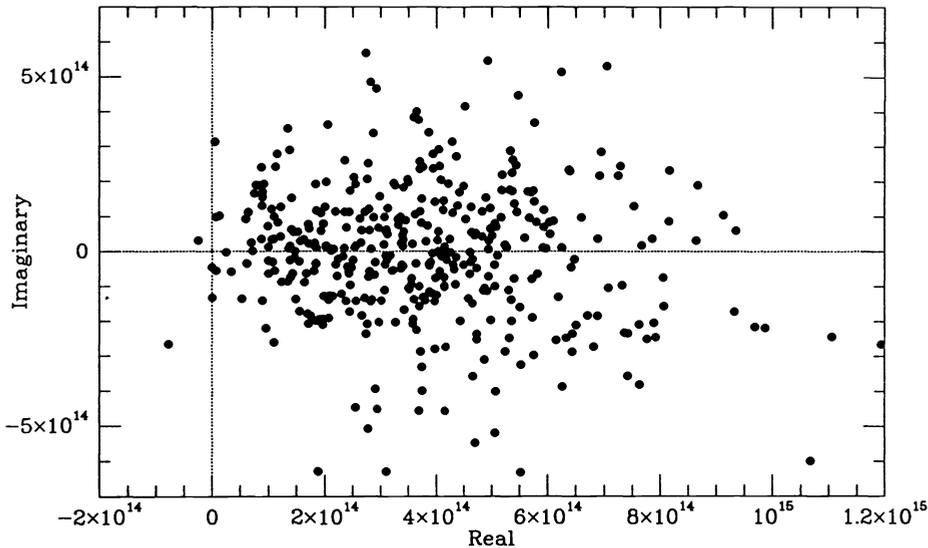
Figure 4. An Argand diagram of individual bispectrum phasors for a stack of 399 calibrator frames, the mean phase of which has SNR = 38.96. The point in bispectrum space being analyzed corresponds to the most redundant triangle in the mask plane. The average closure phase is $-0.75 \pm 1.47°$.

of bispectrum phasors that are not only lower in their mean modulus and SNR, but biased in the positive Φ direction. The mean closure phase is $29.5 \pm 2.46°$. We found that in all such cases where the closure phase is biased, the phase does not appear to vary with time, but, to within the errors, remains consistently at the same bias level for long periods of time.

Upon further investigation, we find that these bias levels not only remain fairly constant with time for a given position of the mask, but exhibit a marked degree of regularity within bispectrum space. To demonstrate this regularity, we consider all the triangles formed by holding the baseline between two apertures constant and allowing the third corner of the triangle to sweep across the other aperture in a raster scan format. Such a procedure is analogous to holding one baseline vector in the uv plane constant and moving a second uv vector across a different region of the power spectrum. The third vector is completely defined by the first two, so it similarly sweeps across a region of the plane. In Figure 7, we plot the closure phases of our calibrator source for two different mask rotations as a function of the location of the second baseline vector in the uv plane.

The darkest contour corresponds to zero closure phase; the others occur at intervals of 5°. The peak in the power spectrum in the displayed region occurs almost exactly at the center of the plots, where the closure phase is unbiased. Assuming the mask was perfectly centered along the optical axis, the two mask positions correspond to the same physical portions of the mirror being covered, assuming the rotations differ exactly by the prescribed 120°. Any differences in the bias may be partly attributable to the fact that the relative degree of rotation was not exactly 120°, or that the mask was not precisely aligned. However, we find the consistency of the gradient in the biases rather remarkable, especially considering we have never encountered consistent biases in FFA data sets taken less than a minute apart, while these two data sets were taken more than forty minutes apart with eleven other observations in between. We also find consistent gradients for the data sets taken between these two. While we do not have an explanation for the cause of the bias at this time, we will at least use its regularity to aid us in the reconstruction of images as described below.
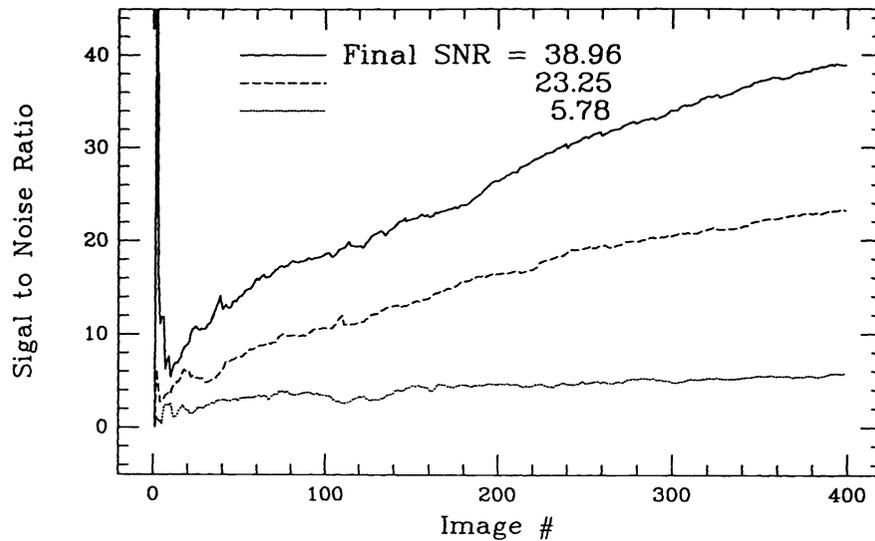
SNR of Mean Bispectrum Phasor

Figure 5. The SNR of the mean bispectrum phasor for three different points within bispectrum space as a function of the number of images used. The solid line corresponds to the equilateral triangle in the mask plane centered within each aperture, while the dashed and dotted lines correspond to the same triangle rotated progressively away from the centers of the holes. Note that in all cases, the SNR appears to increase as the square root of the number of images.
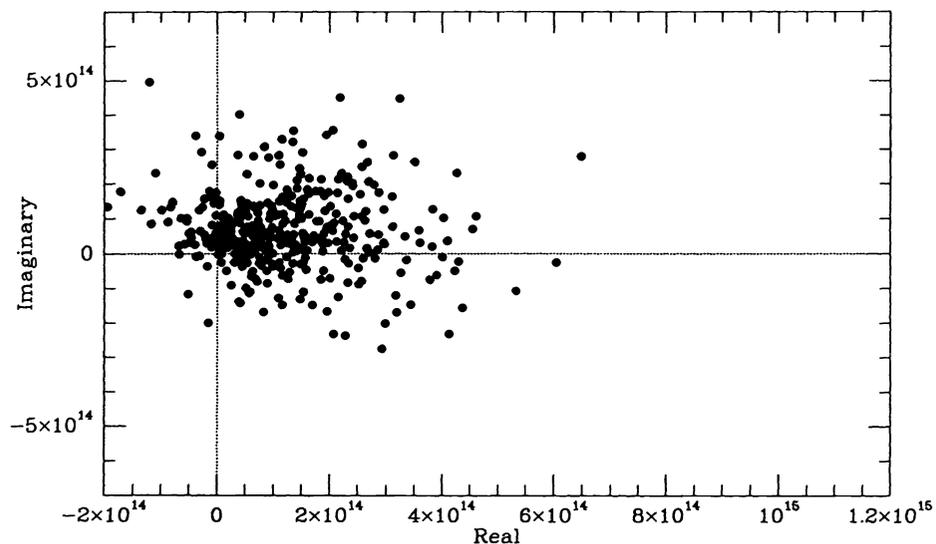
Bispectrum Phasors with S/N=23.25

Figure 6. An Argand diagram of individual bispectrum phasors for a stack of 399 calibrator frames. In contrast to Figure 4, the uv points used do not lie in the peaks of the power spectrum. The result is an ensemble of bispectrum phasors that are not only lower in their mean modulus and SNR, but biased in the positive Φ (clockwise) direction by $29.5 \pm 2.46°$.

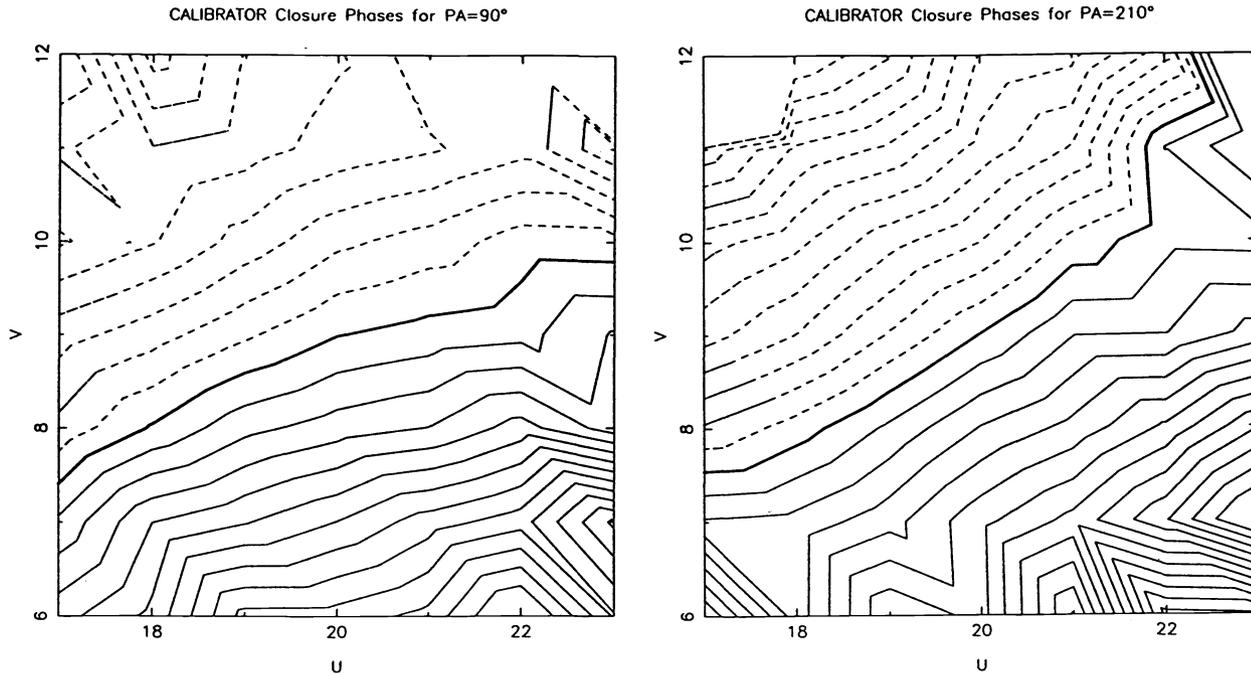CALIBRATOR Closure Phases for PA=90°　　　　　　CALIBRATOR Closure Phases for PA=210°

Figure 7. The closure phases of the calibrator source as a function of the location of one of the baseline vectors in the uv plane. The plots are for two position angles of the mask, differing by a 120° rotation. The darkest contour corresponds to zero closure phase, the lighter and dashed contours occurring at intervals of ±5°.

5. IMAGE RECONSTRUCTION

To reconstruct images from the derived visibilities and closure phases, we use the hybrid mapping package developed at Caltech for the processing of Very Long Baseline Interferometry (VLBI) data. This software implements a version of self-calibration[5,6], whereby the visibilities, closure phases, and errors on each, are combined with an initial model to produce a consistent set of Fourier phases and amplitudes. These Fourier components are subsequently inverted and a dirty beam produced. After CLEANing[7] the dirty map, the resulting model can be fed back at the self-calibration stage, and the three stage process repeated. This iterative approach has been shown to be rather robust and successful for reconstructing images from this type of data, and insensitive to the choice of an initial model[8].

Because of its origins in radio interferometry, it is difficult to use the VLBI software and utilize the full uv coverage provided by a multi-r_o sized mask and/or a broad bandwidth. One must assign the equivalent of a radio telescope station to each aperture, or portion thereof, and then only pure rotations of these station about an arbitrary axis are allowed. For our data, we assign stations corresponding to the centers of each aperture in the three-hole mask. Wishing to take advantage of the breadth of uv coverage provided by even a single pair of apertures, we then allow three "sub-rotations" of the stations in the clockwise and counter-clockwise direction about the nominal position, in steps of two degrees each. We thereby sample the full breadth of coverage of the mask in the uv plane, but only in the tangential direction. Figure 8a is a plot of the uv coverage for the two position angles (30° apart) of the mask we use. In total, we sample 42 points in the half-plane.

As demonstrated in section 4, the closure phase of the calibrator is zero only for the triangle with vertices at the center of each aperture. However, because the bias appears to be relatively constant over many data sets, we decide to try and calibrate the object closure phases by subtracting the closure phases of the calibrator at each of the measured "sub-rotations" of the triangle.

The results of supplying the VLBI software with these calibrated closure phases and visibilities are illustrated in
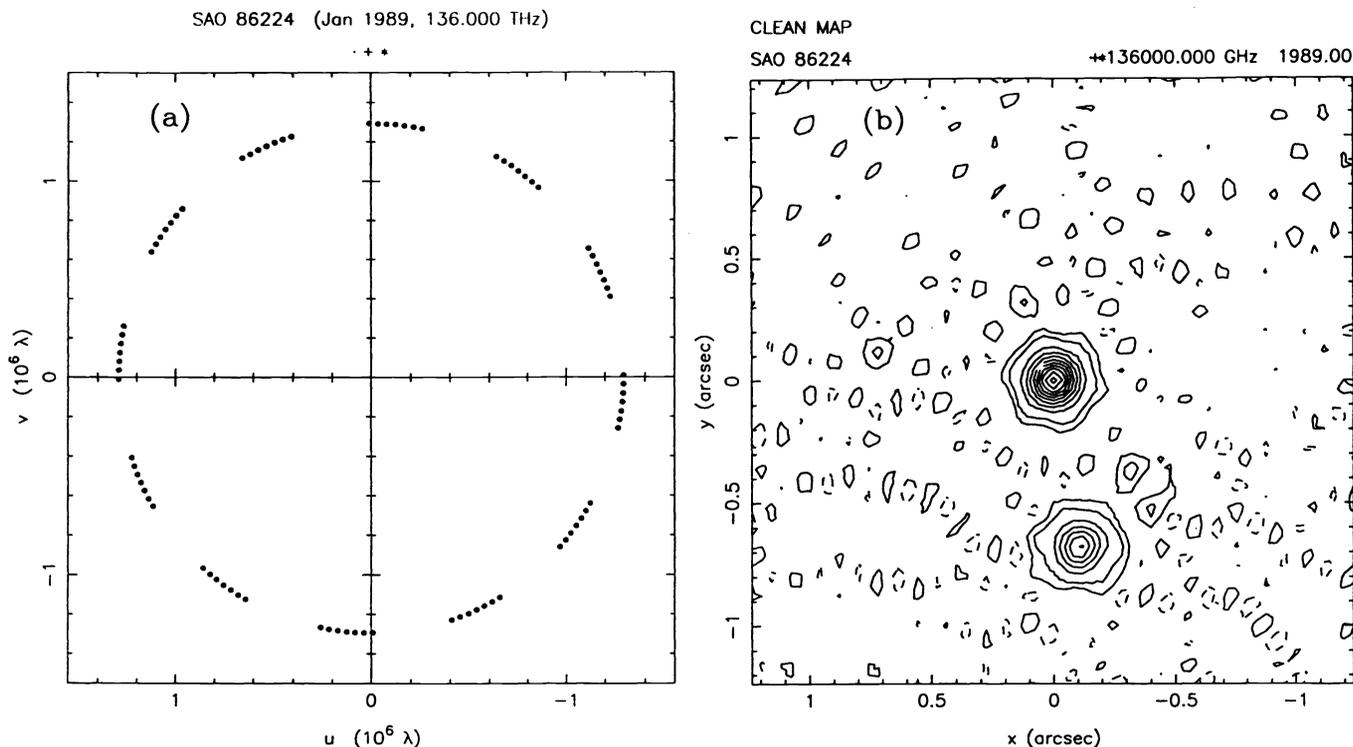
SAO 86224 (Jan 1989, 136.000 THz)

CLEAN MAP

SAO 86224   ++136000.000 GHz 1989.00

Figure 8. (a) The total uv coverage for two position angles of the mask and each "sub-rotation" therein. This represents our coverage for the reconstructed image (b) of SAO 86224 produced after nine iterations of the self-calibration cycle. The contours are at -2, 2, 5, 15, 25, 35, 45, 55, 65, 75, 85, and 95% of the peak.

Figure 8b, which shows the Clean map, or reconstructed image, resulting from nine iterations of the self-calibration cycle. The dynamic range of the image is reasonably good, approaching 20:1, and the binary nature of the object is clearly distinguishable. Figure 9 shows a comparison of the measured visibilities and closure phases with the visibilities and phases produced by the principal components that comprise reconstructed image.

We also used a least squares model fitting program, also part of the VLBI package, to find the best fitting binary model of the data. We provided as input the two strongest components in the Clean map, and then allowed the model fitting program to vary the fluxes, radial distance and position angle of the two sources. The resulting model is almost identical to that produced by self-calibration, and as Figure 10 shows, the visibilities are only slightly better in the model produced by the model fitting program. The closure phases of this model, however, match the measured closure phases significantly better than do the the closure phases of the reconstructed image.

We next tried running self-calibration on the same data but with uncalibrated closure phases. We found we do get an image of the binary, but of noticeably lower dynamic range than the image produced previously. In Figure 11 we plot the closure phases of the best fitting model to the uncalibrated data. The visibilities and closure phases of the model are virtually identical to those derived before, but the measured (uncalibrated) closure phases do not match the model closure phases very well at all. The fact that both initial data sets produce virtually the same model image visibilities and closure phases, while only the calibrated closure phases match the model reasonably well, suggests that phase calibration is warranted and even necessary for producing images of the highest dynamic range.
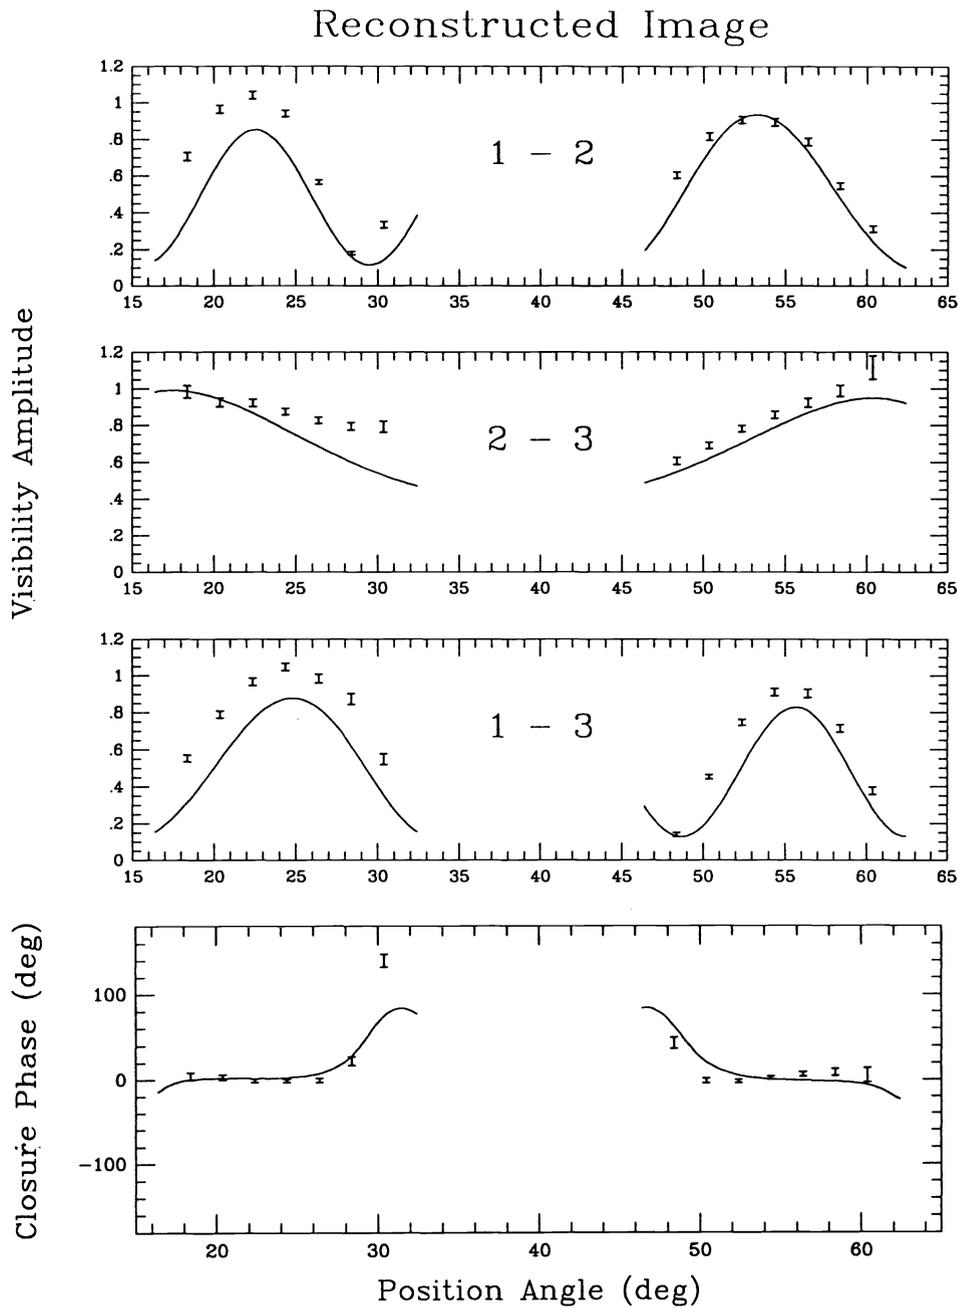
Figure 9. A comparison of the visibilities and closure phases for the measured data versus the reconstructed image. The error bars represent one sigma ranges for the data. The solid line represents the values for the reconstructed image.
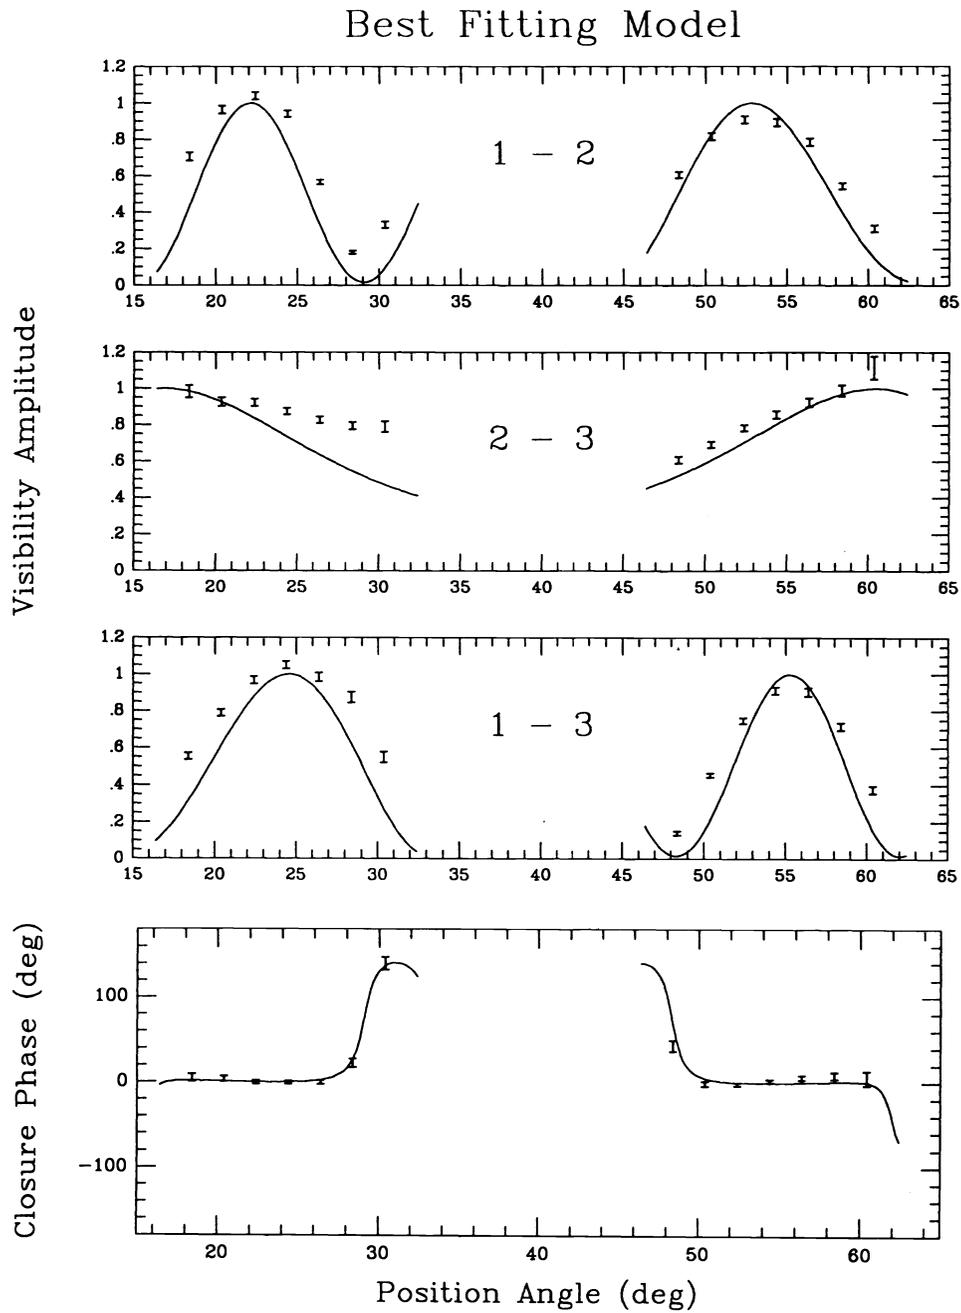
Best Fitting Model

Figure 10. A comparison of the visibilities and closure phases for the measured data versus the best fitting model image. Again, the error bars represent one sigma ranges for the data while the solid line represents the values for the model. Note the improved fit to the closure phases.
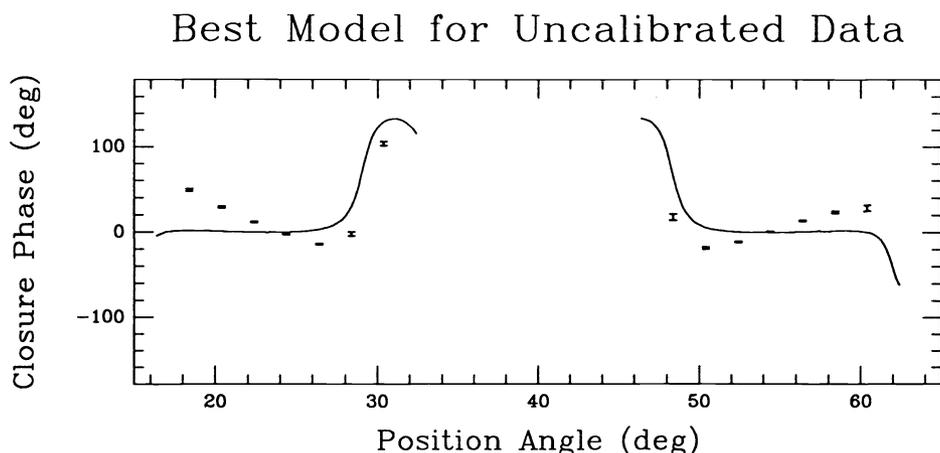
Best Model for Uncalibrated Data

Figure 11. A comparison of the closure phases for the measured, uncalibrated data versus the model which best fits these data. Note that the model closure phases are indistinguishable from those of the previous model, but that they significantly differ from the uncalibrated data.

6. CONCLUSION

Our preliminary investigation of NRM imaging at infrared wavelengths demonstrates the viability of such techniques for producing high resolution images, as well as for serving as a useful tool for investigating general aspects of imaging. We find that the noise properties of our system are in general well understood, with no systematic problems affecting the SNR of either the power spectrum or the bispectrum. While the bias in the power spectrum is well understood, we still have much work to do to fully understand the bias evident in the bispectrum of our data. We are currently preparing various tests and analyses designed specifically for this purpose.

We have also demonstrated that with the use of large, even multi-r_o, apertures and a large bandwidth, it is possible to recover visibility and phase in a significantly large fraction of the uv plane with a single mask at just a single position angle. We still must determine at what point as one increases the size of the apertures or the bandwidth, the decrease in visibility and irregularity in the closure phase bias outweigh the benefits of the increased uv coverage.

7. ACKNOWLEDGMENTS

8. REFERENCES

1. P.W.Gorham, A.M.Ghez, S.R.Kulkarni, T.Nakajima, G.Neugebauer, J.B.Oke, T.A.Prince, "Diffraction-limited imaging. III. 30mas closure phase imaging of six binary stars with the Hale 5m telescope," *Astron. J.,* Vol 98, pp. 1783-1799, 1989.
2. T.Nakajima, S.R.Kulkarni, P.W.Gorham, A.M.Ghez, G.Neugebauer, J.B.Oke, T.A.Prince, A.C.S.Readhead, "Diffraction-limited imaging. II. Optical aperture-synthesis images of two binary stars," *Astron. J.,* Vol 97, pp. 1510-1521,1989.
3. C.A.Haniff, C.D.Mackay, D.J.Titterington, D.Sivia, J.E.Baldwin, P.J.Warner, "The first images from optical aperture synthesis," *Nature,* Vol 328, pp. 694-696, 1987.
4. G.R.Ayers, "Correlation techniques for imaging through atmospheric turbulence," Ph.D. Thesis, Imperial

College of Science, Technology and Medicine, University of London, p. 105, 1988.

5. T.J.Cornwell, P.J.Wilkinson, "A new method for making maps with unstable radio interferometers," *Mon. Not. R. astr. Soc.*, Vol 196, pp. 1067-1086, 1981.

6. T.J.Pearson, A.C.S.Readhead, "Image formation by self-calibration in radio astronomy," *Ann. Rev. Astron. Astrophys.*, Vol 22, pp. 197-130, 1984.

7. J.A.Högbom, "Aperture synthesis with a non-regular distribution of interferometer baselines," *Astron. Astrophys. Suppl.*, Vol 15, pp. 417-426, 1974.

8. C.A.Haniff, personal communication, 1990.