Diffraction-limited imaging with partially redundant masks. I. Infrared imaging of bright objects

Christopher A. Haniff*

Division of Physics, Mathematics and Astronomy, MS 105-24, California Institute of Technology,
Pasadena, California 91125

David F. Buscher

Naval Research Laboratory/U.S. Naval Observatory Optical Interferometer Project, U.S. Naval Observatory, AD5,
34th Street and Massachusetts Avenue, N.W., Washington, D.C. 20392

The utility of partially redundant pupil geometries has been studied in the context of near-infrared speckle imaging with ground-based telescopes. Using both numerical simulations and experimental data collected with a 4-m-class telescope, we find that the decrease in redundancy resulting from apodizing the telescope pupil results in an enhancement of the quality of reconstructed images at high light levels. This improvement in imaging fidelity is particularly valuable when short-term variations in the statistics of the atmosphere make the seeing calibration of speckle interferograms difficult. However, the use of an apodizing mask necessarily restricts the faintest source that can be imaged, leading to a loss in sensitivity of one to two magnitudes. For many of the brighter near-infrared astrophysical sources in the sky that have been the subject of previous speckle-imaging studies, the use of a partially redundant pupil is expected to enhance the fidelity of the imaging procedure considerably.

1. INTRODUCTION

The recent advent of two-dimensional near-infrared-array detectors[1,2] has stimulated considerable effort among the astronomical community in the application of methods for speckle imaging at infrared wavelengths. In contrast to the situation in optics in which most of the development of advanced speckle-imaging techniques has been performed, existing near-infrared array detectors are unable to count and to time tag individual photon events. These detectors operate in a mode that is analogous to that of a shuttered CCD, albeit with much poorer noise characteristics and with the added feature of an increasing thermal background beyond approximately 2.5 μm. These problems, together with the relatively low duty cycle of such detectors when they are compared with optical photon-counting cameras, have meant that speckle observations of the faintest infrared sources are currently not practicable. As a result of this, researchers have tended to concentrate on the detailed study of a small number of bright prototypical targets.[3,4]

The majority of theoretical studies of speckle imaging[5-7] have restricted their attention to the photon-starved system in an attempt to ascertain the fundamental limitations of the various image-reconstruction techniques available. On the other hand, existing near-infrared speckle cameras operate in a high-light-level regime, in which the most important contributions to the noise on the interferometric observable being measured arise not from the statistical fluctuations in the detected photon counts but from the redundancy noise associated with the telescope pupil geometry. Under these conditions, the flexibility afforded by a variable pupil geometry can be-

come increasingly important in optimizing any particular imaging experiment. Despite the possible importance of this role, speckle-imaging experiments with large ground-based telescopes have relied almost exclusively on unobscured pupil geometries, although the rationale for this choice is rarely justified. To date the only different approach that has been examined in depth is the nonredundant masking scheme.[8,9] In this method the telescope pupil is covered with an opaque screen that transmits light only through a number of nonredundantly arranged subpupils, each of which is smaller than the characteristic seeing scale size, r_0. Although such nonredundant pupils often collect only a small fraction of the light incident upon the primary mirror, they facilitate high-dynamic-range imaging and are well suited to brighter astronomical targets. Results obtained with this method at both optical[10-13] and infrared[14] parts of the spectrum have been presented by a number of groups.

A related hybrid approach, in which a partially redundant apodizing mask is employed, was recently suggested by Haniff *et al.*,[15] who advocated the use of an annular pupil with a width comparable to r_0. In the present paper we examine the rationale for this approach, assuming atmospheric and detector characteristics appropriate for the near-infrared imaging of relatively bright sources. Under these conditions the use of a partially redundant pupil may be favored for two main reasons: first because it can reduce the effects of atmospheric noise and second because it can provide improved resistance to the seeing variations that affect the fidelity of the image-reconstruction procedure. The second of these issues has, to our knowledge, never received any attention in the literature, despite the fact that it is generally agreed that the seeing miscali-

bration between observations of an object and its unresolved reference source represents the most important practical limitation to the speckle technique.[16,17] We support our arguments with both numerical simulations and real data obtained with a near-infrared speckle camera. These reveal the versatility of the use of partially redundant pupil geometries and demonstrate that such methods are potentially valuable for a range of astrophysically interesting sources.

In Section 2 we develop the parallel themes of pupil redundancy, atmospheric noise, and seeing calibration, and we calculate the dependencies of the latter two quantities on the redundancy of an annular pupil of variable width. In Section 3 we describe the procedures we have developed to reduce interferometric data obtained with such annular pupils, and in Section 4 we present image reconstructions from simulated data to assess the relevance of the use of partially redundant pupils in a practical imaging environment. Our conclusions are given in Section 5.

2. PUPIL REDUNDANCY, ATMOSPHERIC NOISE, AND SEEING CALIBRATION

In considering the consequences of the pupil geometry, it is instructive to consider first two extreme cases, namely, a fully filled (i.e., unapodized) aperture and one that has been masked off apart from a nonredundant array of small holes. These provide illustrative examples of the effects of the pupil shape on redundancy noise and seeing calibration at high light levels and on the signal-to-noise ratios (SNR's) of power spectrum and bispectrum measurements made at lower light levels, and it is hoped that they will clarify the subsequent treatment of partially redundant pupils presented in Subsection 2.B.

A. Comparison of Fully Filled and Nonredundant Pupils: Qualitative Results

1. Atmospheric Noise
A convenient framework in which to examine the effects of pupil redundancy is one in which the telescope aperture is envisaged as being composed of a large number of subpupils, each of diameter $\leq r_0$. The distortions that are due to atmospheric turbulence are assumed to be constant in phase across each subpupil but to vary randomly between subpupils and between exposures. For excellent introductions to such an heuristic description the reader is referred to the reviews of Roddier[18] and Goodman.[19] Any pair of the subpupils can be considered a Michelson interferometer contributing to a single spatial frequency in the image. The spatial frequency is defined by the vector separation of the subpupils. If any particular baseline is repeated in the pupil, i.e., if more than one pair of subpupils in an aperture has the same vector separation, the pupil is said to be *redundant*, and the redundancy of the baseline is given by the number of times it is repeated.

In the simplest nonredundant configuration, one can imagine the pupil being apodized so that only three appropriately located subpupils transmit light to the detector. A single-speckle interferogram of an unresolved source will then show three coincident Airy patterns over which three sets of fringes are superposed. Individual specklegrams obtained through separate realizations of the atmosphere will differ only in the fringe phases, i.e., the

locations of the fringe maxima with respect to some fiducial reference point on the detector. Consequently the measured values of both the power spectrum and bispectrum will be unaffected by the instantaneous fluctuations of the fringe phases. Sequential estimates of these two "good" observables will therefore vary only because of the effects of readout and photon noise.

On the other hand, if an unobstructed pupil is employed or if the pupil is apodized so that a redundant configuration of subapertures remains, the power at any single point in the Fourier plane will be the incoherent sum of contributions from many pairs of subpupils. Consequently the measured values of the power spectrum and bispectrum will fluctuate about their mean values as separate specklegrams sample different realizations of the atmospheric wave-front perturbations. This additional variation, commonly referred to as atmospheric or redundancy noise, limits the instantaneous SNR of the power spectrum and bispectrum to unity at high light levels, and so it should not be surprising that reducing the
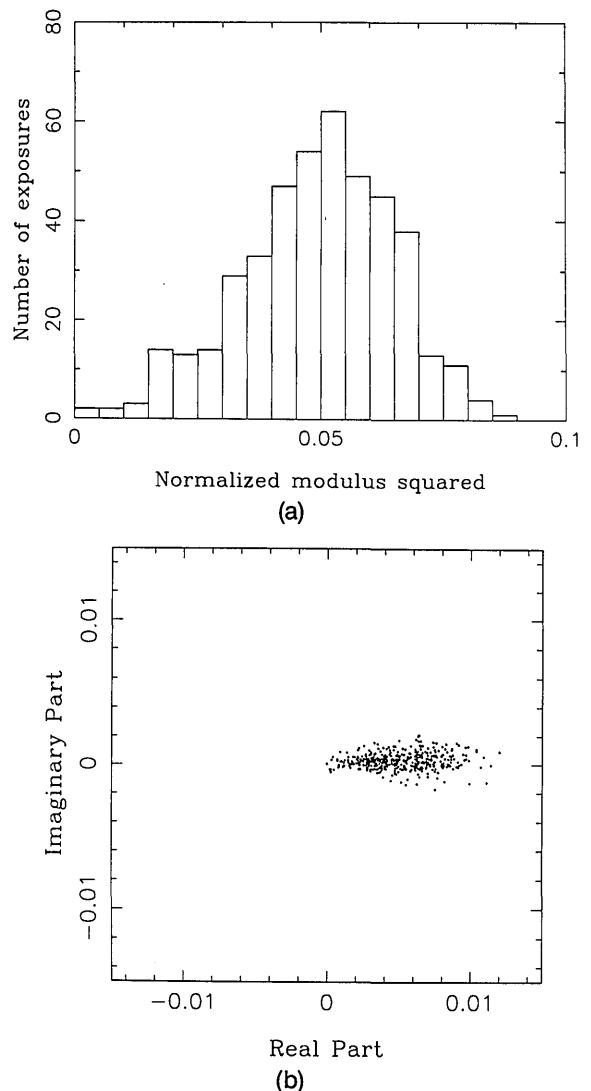

Fig. 1. (a) Histogram of 485 sequential measurements of the power spectrum magnitude at a fixed spatial frequency corresponding to one pair of subpupils in a nonredundantly masked pupil. (b) Corresponding Argand diagram for measurements of a single point in the bispectrum. See the text for details of these point-source observations.
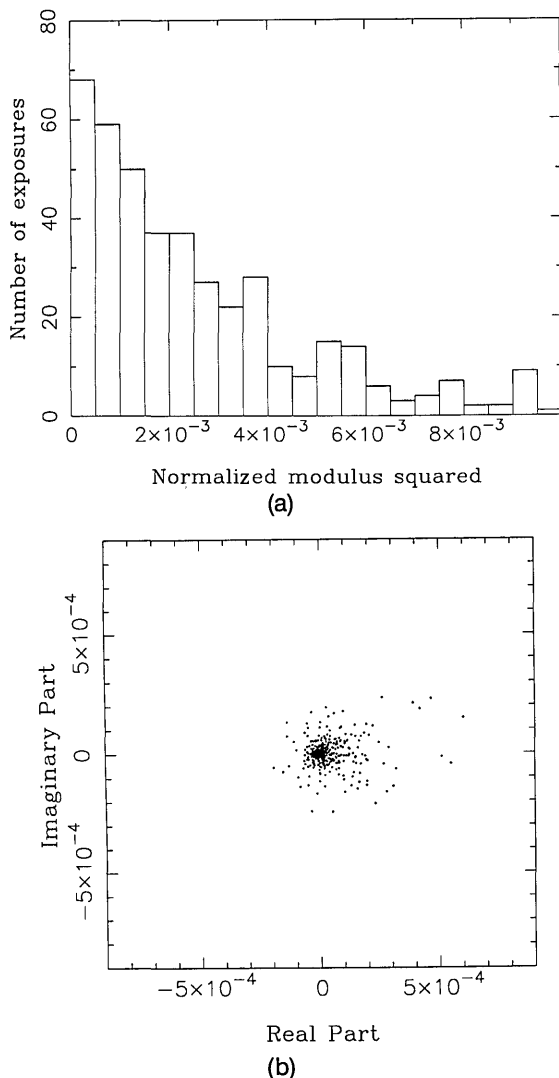
Fig. 2. (a) Power spectrum magnitude and (b) bispectrum data for an observation of an unresolved source obtained through an unapodized pupil. The differences between these data and those displayed in Fig. 1 are primarily a result of the increased redundancy noise owing to the unapodized pupil geometry.

redundancy of the telescope pupil may be advantageous if the noise attributable to the pupil geometry dominates over other sources of noise.

These effects are illustrated by the experimental data in Figs. 1 and 2, which were recorded with the KPNO speckle camera[3] on the Mayall 4-m telescope in May 1989. We have presented the results in the form of histograms of measured power spectrum values and scatter diagrams of bispectrum values for some *fixed* spatial-frequency coordinate(s) over an ensemble of speckle interferograms. Figure 1 shows typical distributions obtained for a bright source observed with a nonredundantly masked pupil. The pupil mask used comprised four nonredundantly arranged pupils, each of diameter 40 cm. Observations were made of the bright star α Bootis at a wavelength of 2.2 μm ($\Delta\lambda = 0.09$ μm); an exposure time of 50 ms and a pixel scale of 0.058″ were used. Figure 1(a) shows a histogram of the measured values of the source power spectrum modulus (i.e., the visibility amplitude squared) at a baseline of 1.54 m, while Fig. 1(b) shows an Argand diagram of the measured bispectrum values for a closing loop

of baselines of lengths 0.70, 1.54, and 1.85 m. For this particular source, the effects of photon and readout noise are very small and the observed statistical variations in the measurements from specklegram to specklegram reflect the finite subpupil size and exposure time, which are nonnegligible in relation to the spatial- and temporal-coherence scales of the atmosphere (r_0 and t_0). The small corrugations present in the optical wave fronts across individual subpupils and the fluctuations of the wave-front phase during separate short exposures lead to random variations in the fringe parameters that are not taken into account by the simple model described above. Because the correlations in the perturbed wave front do in fact vary smoothly with distance in the pupil, the simple definition of redundancy derived from the discrete heuristic model should be understood in light of the fact that the atmospheric noise increases continuously as a function of the pupil area.

The observed power spectrum and bispectrum estimates are consistent with those expected for an unresolved source, showing a Gaussian-shaped power spectrum distribution centered at a normalized source visibility of 0.89 and a tightly clustered bispectrum-phase distribution with a mean value of 3.0° ± 0.3°. The small but significant discrepancy between the mean bispectrum phase and the expected value of zero can be attributed to fixed telescope-based aberrations that we have found to be stable over time scales of several hours and that are probably due to slight imperfections in the primary-mirror support system. An interesting feature of Fig. 1(b) is the strong elongation of the bispectrum-phasor distribution along the real axis. Although most of the bispectrum-phase measurements are close to the expected value of zero degrees, the range in the bispectral modulus is more than a factor of 10, a fact that highlights the need to parameterize the phase and amplitude variations of the bispectrum separately. Because only the bispectrum phase and not its modulus is routinely required for image recovery, it is preferable to characterize the statistical variability of the bispectrum measurements in terms of a phase-error parameter that quantifies the variability of the argument of the measured bispectrum but that does not reflect its changing amplitude. Throughout this paper we use the term phase error to denote such a quantity and follow the definition of Buscher[20]:

$$\sigma_\phi \equiv \sigma_\perp / |\langle S \rangle|, \qquad (1)$$

where σ_ϕ is the phase error in radians, $\sigma_\perp{}^2$ is the variance of the bispectrum phasor in the direction perpendicular to the mean phasor, and $\langle S \rangle$ is the mean phasor.

Figure 2 shows data similar to those presented in Fig. 1, obtained with the same telescope, detector, filter, and exposure time and under similar seeing conditions (~0.9″ at 2.2 μm) but without the use of a pupil-plane mask. Because of an alignment problem with the speckle camera on the night of the observations, the effective outer diameter of the telescope pupil was only 3.2 m for these data, which were obtained for the unresolved source 31 Leonis (HR 3980) on May 22, 1989. As for the data of Fig. 1, the effects of photon and readout noise on the measurements are unimportant, but, in contrast to those data, the statistical fluctuations of the individual measurements are dominated by the redundancy noise arising from the pupil

geometry and not, for example, by the finite exposure time. It is at once evident that the distributions of the measured quantities have been considerably broadened as a result of the pupil configuration. The power spectrum measurements do follow the predicted negative-exponential distribution, and thus a single measurement has a unit SNR. Although the mean bispectrum phase ($3.5° \pm 5°$) is close to the expected value of zero degrees, the error on the mean value is still high despite the large number of independent measurements in this data set. The difference in the shapes of the bispectrum distributions is perhaps the clearest sign of the redundancy noise and leaves no doubt as to the utility of nonredundant pupil geometries if the source being observed is sufficiently bright.

2. Visibility Calibration

The results of Subsection 2.A.1 confirm that pupil redundancy can strongly affect the random noise on the measurement of the power spectrum and bispectrum at high light levels. However, there is another corollary of the pupil geometry, which is of more concern at both high and low light levels because it is a systematic effect in the sense that it does not change randomly between exposures and thus cannot simply be removed by averaging many exposures; this is the problem of amplitude calibration.

Although there exist algorithms for the blind recovery of diffraction-limited images from specklegrams,[21,22] the most frequently used procedures[23] all require observations of an unresolved source to measure the statistical reduction in visibility amplitude that are due to the combined response of the atmosphere and the optics. The magnitude of this reduction, especially the component that is due to the atmosphere, varies with time and with position in the sky, and so observations of the source and the calibrator must be made as close together both in time and in space as is possible. Even if great care is taken, however, residual differences in the seeing between the source and the calibrator are likely to remain. It is important to note that the averaging of many exposures will not reduce a seeing mismatch as effectively as it can reduce the effects of atmospheric and detector-related noise. This is because, whereas the latter are effectively uncorrelated between exposures, variations in the seeing occur on time scales ranging from seconds to hours. While rapid switching between the source and calibrator stars will remove the effects of the longer-term seeing variations, it will not be so effective against fluctuations that occur on time scales that are comparable with the switching. Such switching will also be unable to remove differences in seeing that are dependent on the positions in the sky of the source and the calibrator, for example, zenith-angle dependencies and seeing phenomena that are localized with respect to the telescope.

An additional problem arises because the influence of a change in the mean seeing will be correlated in the spatial-frequency domain, tending to depress or raise the mean amplitudes uniformly in the power spectrum. This is in contrast to the more-random effects of detector and atmospheric noise, which are likely to be uncorrelated between widely different spatial frequencies. A small seeing miscalibration may therefore have a much more significant effect on the reconstructed image than a similar amount of atmospheric or detector noise, because the latter is

uncorrelated between exposures and between spatial frequencies.

In light of the above discussion, a pertinent question is "To what extent is miscalibration a function of the pupil geometry?" If we consider a nonredundant mask with holes much smaller than r_0, it is straightforward to argue that the visibility of the fringes produced will be relatively unaffected by an increase in the seeing-cell size, because it will be nearly unity to start with and cannot be increased any further. Conversely, in the case of an unapodized circular aperture the power spectrum amplitude will be greatly increased if the seeing improves, because the size of the patches in the pupil that contribute coherently to the power spectrum is determined by the atmosphere and *not* by the pupil. Korff[24] has shown that, for midrange spatial frequencies, the mean squared modulation-transfer-function of a pupil that is large compared with r_0 will vary as r_0^2. At spatial frequencies comparable with r_0/λ, the effect of a change in seeing will be even greater, because the width of the seeing spike caused by short-range correlations in the wave front is approximately r_0/λ and the modulus of the speckle-transfer function is many times larger within this spike than in the plateau region at higher spatial frequencies. This will be a particularly important side effect at near-infrared wavelengths at which the ratio of the pupil diameter, D, to r_0 can be as small as 5 for a moderately sized telescope. We can conclude, then, that a nonredundant geometry, at least in the case in which the subapertures are smaller than r_0, will be much less sensitive to a change in the mean seeing than a redundant pupil configuration.

3. Limiting Magnitudes

Understandably, the use of nonredundant pupils comprising subelements smaller than r_0 has been criticized because it is wasteful of the total flux collected by the unapodized pupil. Indeed, it is at low light levels that pupil redundancy becomes advantageous. Using the heuristic model for the fully filled aperture introduced above, we can see that the power at a given spatial frequency in the power spectrum is the sum of many independent contributions from different redundant baselines. Although this power may fluctuate widely, its mean value will be greater than that measured with a single pair of subapertures. Similarly, the signal at a point in the bispectrum will contain contributions from many triplets of closing baselines in the pupil and on average will have a greater amplitude than a measurement made with a nonredundant pupil. In the photon-noise-limited system, this increase in signal power afforded by pupil redundancy may be offset by a corresponding increase in the noise level, because the latter is set by the total number of photons collected, which will be much greater in the case of an unapodized pupil than in the case of a nonredundant one. This consideration will be dealt with in a later paper but is not applicable to near-infrared speckle interferometry because at both low and high light levels the photon-related noise is insignificant compared with noise from other sources: At low photon rates the readout noise of the detector dominates, whereas at high photon rates the atmospheric noise exceeds the noise that is due to the readout electronics. Thus when detector noise is dominant the noise level becomes independent of the signal

level, and so an increase in the mean signal level implies a concomitant increase in the SNR.

In this sense, therefore, the use of anything but the full unapodized pupil should be expected to compromise the limiting sensitivity of infrared speckle observations to some degree. However, the increase in the mean signal level is a comparatively weak function of the pupil area and hence of redundancy when observation takes place through atmospheric turbulence. If wave-front-phase perturbations are unimportant, the power spectrum and bispectrum signal levels scale as the second and third powers of the pupil redundancy, respectively, but in the presence of random-phase perturbations both quantities have only a linear dependence[22] on pupil redundancy, implying that the loss in sensitivity that is due to pupil apodization need not be severe.

B. Partially Redundant Pupil Geometries: Quantitative Results

1. Rationale and Numerical Procedure

We have seen that the use of a nonredundant pupil can dramatically reduce the noise and miscalibration effects for high-light-level interferometric measurements. Unfortunately, at infrared wavelengths the use of truly nonredundant pupil geometries is of limited utility for a number of reasons. For enough light to be collected to overcome the readout noise of existing detectors, the subaperture diameters must be of the order of the size of the seeing cells (i.e., r_0), and yet, for a truly nonredundant array, the subapertures' center-to-center separation must be at least twice their diameter. When the ratio of the telescope diameter to the seeing-cell size is small, as is the case at infrared wavelengths, this restricts the number of subapertures that can be used and hence the instantaneous Fourier-plane coverage attainable. Furthermore, because of sidelobe effects, the interferometric field of view is restricted to $\sim\lambda/d$, where d is the minimum aperture spacing, so that the mapping of wide fields requires ever more closely spaced and hence smaller subpupils. While this is not usually a problem at optical wavelengths because the objects of interest are much smaller than the seeing limit (i.e., λ/r_0), the infrared sky contains many sources (e.g., evolved stars and protostellar objects) that simultaneously display interesting structures at angular scales comparable with and smaller than the seeing limit. This places severe constraints on the possible nonredundant arrangements of subpupils that are useful, especially if the subpupil dimensions are large.

Another consideration is that, in the readout-noise-limited system, nonredundant pupils are suitable only for the brightest sources. Even though a large fraction of the astrophysically interesting near-infrared sources that can be resolved by interferometric methods are bright, there are other classes of sources, for example, low-mass binary systems, that are much fainter. In such cases the effects of detector noise become important, and so the baseline redundancy offered by unapodized pupil geometries can be valuable in terms of increasing the mean signal level detected, albeit only slowly as the pupil redundancy rises.

Our discussion thus far suggests that there is scope for a method that can exploit both the complete Fourier-plane

coverage and high signal levels provided by fully filled aperture techniques and the resistance to seeing variations that nonredundant schemes afford. One such approach is to apodize the pupil in such a way as to achieve full Fourier-plane coverage while minimizing the pupil redundancy. Hereafter we call this approach partially redundant masking, or PRM. Such a technique would be expected to exhibit properties intermediate in nature between those characteristic of fully redundant and nonredundant pupil geometries and to give rise to a moderate amount of atmospheric noise on bright sources but offer the possibility of imaging sources that are much fainter than those that can be mapped by using the nonredundant masking method.

To quantify the gains that such an hybrid method might achieve and to compare the imaging fidelity attainable with that expected from a fully filled aperture, we have numerically simulated the process of imaging through atmospheric turbulence. In the simulations, random phase screens with a spectrum of fluctuations corresponding to Kolmogorov turbulence were generated by using the method outlined by Buscher.[20] For each phase screen generated, the screen was multiplied by a user-specified pupil function, and the image of a point source seen through the combined screen–pupil function was calculated. Because we are interested in the high-light-level behavior of these systems, no detector or photon noise was added to the simulated interferograms. The images were Fourier transformed, and the power spectrum and bispectrum were calculated on an image-by-image basis. The ensemble averages and the variances of these quantities were accumulated over a large number of interferograms (typically 10^4), and these were then used to derive the SNR of the power spectrum and the phase error of the bispectrum for a number of pupil configurations.

We have applied these simulations to two partially redundant pupil geometries and compared the results with those computed for a filled pupil. Although a large range of partially redundant pupil geometries may be advantageous, we have chosen to examine only annular configurations because they have a number of immediately
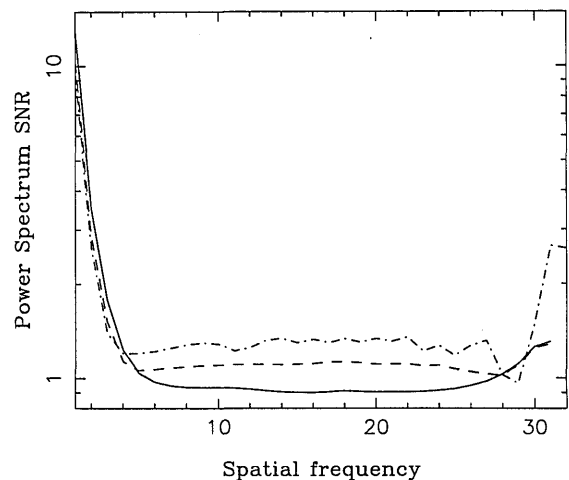
Fig. 3. Noise resulting from atmospherically induced fluctuations in the high-light-level-image power spectrum for the three different pupil geometries discussed in the text. The solid curve shows the SNR of the power spectrum for the unapodized pupil, the dashed curve shows that for the r_0-thick annulus and the dotted–dashed curve shows that for the 0.5-r_0-thick annulus.
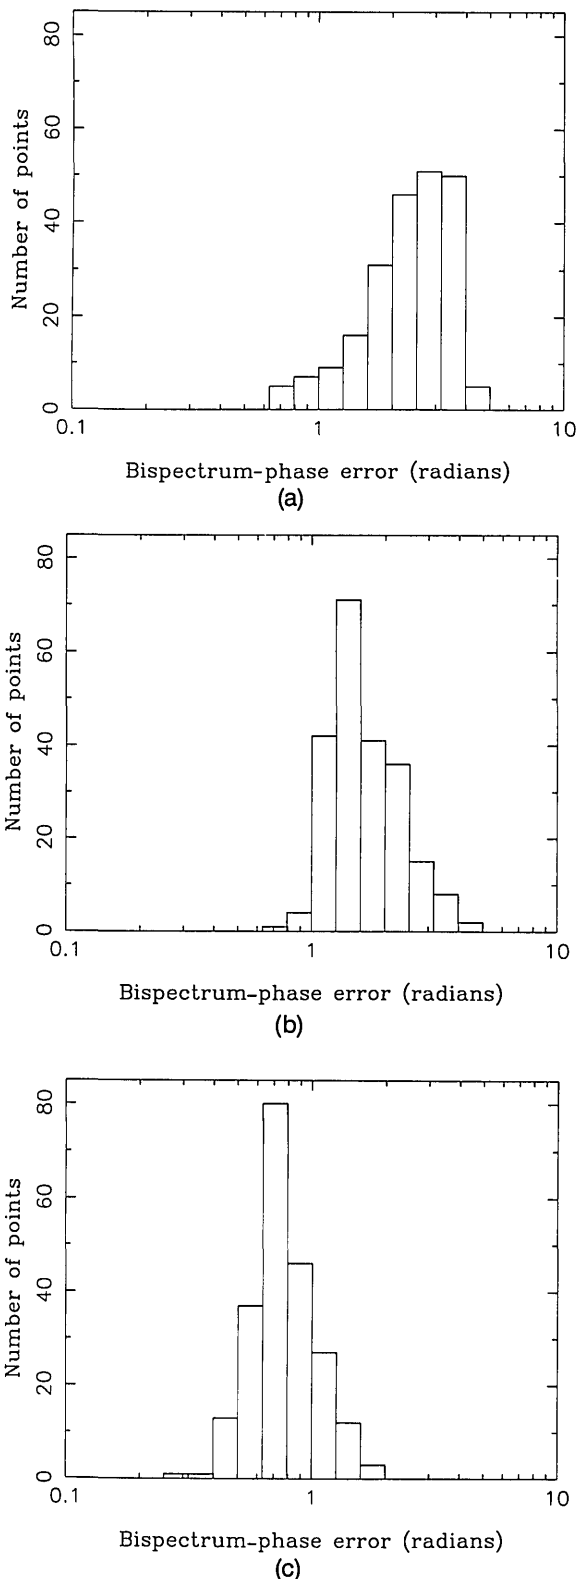
Fig. 4. Histograms of the high-light-level phase errors for a sample set of 220 bispectrum points (a) for the unapodized aperture, (b) for the r_0-thick annulus, and (c) for the 0.5-r_0-thick annulus.

recognizable properties, which include complete spatial-frequency coverage coupled with low redundancy. For these simulations a ratio of D, the outer pupil diameter, to r_0 of 8, which is typical of large (4-m-class) telescopes operating at 2.2 μm, was assumed. We compared an un-

apodized aperture with no central obscuration—this was found to be unimportant if it was less than half the diameter of the pupil—with two annular apertures of widths r_0 and $r_0/2$. The outer diameters of the two annuli were $8r_0$ and $7r_0$, the smaller outer diameter of the thinner annulus being chosen so that the high-frequency peaks of the power-spectrum-transfer functions of the two annuli occurred at approximately the same spatial-frequency radius.

2. High-Light-Level SNR Considerations

The main results of our high-light-level numerical simulations are displayed in Figs. 3 and 4. We consider first the repercussions of the pupil geometry on the SNR of power spectrum measurements. In Fig. 3 we show radial slices through the power spectrum SNR surfaces for our three adopted pupil functions. The behavior of the curve of the fully filled aperture is consistent with theoretical expectations and can be divided into three main regions: a low-frequency spike where small-scale correlations in the wave-front perturbations enhance the coherent addition of signal phasors, a plateau where the SNR saturates at unity, and a rising tail at the highest spatial frequencies where the pupil becomes effectively nonredundant. Both the PRM curves show the same basic form; they differ only in the absolute level of the plateau region, which increases slightly as the annulus width is reduced. Interestingly, the saturation of the SNR around unity that one expects from an addition of random phasors occurs even when there is only a small amount of redundancy. We conclude that there is little to be gained in terms of reducing the measurement noise on the visibility amplitudes when a partially redundant mask is used in this system.

Turning now to the noise on the bispectrum measurements, we have summarized our results in Fig. 4. Rather than presenting SNR curves for a typical slice through the four-dimensional bispectrum, we have chosen to display our results in the form of histograms of the bispectrum-phase errors measured for a sample of 220 bispectral points distributed randomly throughout the full nondegenerate bispectrum hypervolume. The spread in the histograms is due to the real variation of the phase error as a function of position in the bispectral space rather than to residual inaccuracies in the ensemble averages used to compute the phase errors.

Our data show clearly that the phase errors vary significantly for the different pupil geometries, shifting to lower values as the annular width is reduced. As a measure of the typical values of the phase-error distributions, we have calculated their medians, since the value of the median is less sensitive to outliers in the distribution than is the value of the mean. More importantly, the median can also be interpreted in relation to the pragmatic decisions involved during Fourier-phase retrieval. Because the number of independent bispectrum phases is usually many orders of magnitude larger than the number of Fourier phases, a large fraction of the bispectrum data can be discarded, for instance, to save computation time, without seriously affecting the quality of a final image reconstruction. The median can be thought of in this context as the remaining worst-case error if the least reliable half of the data is ignored. The median phase errors for the three pupil geometries are, in descending order of
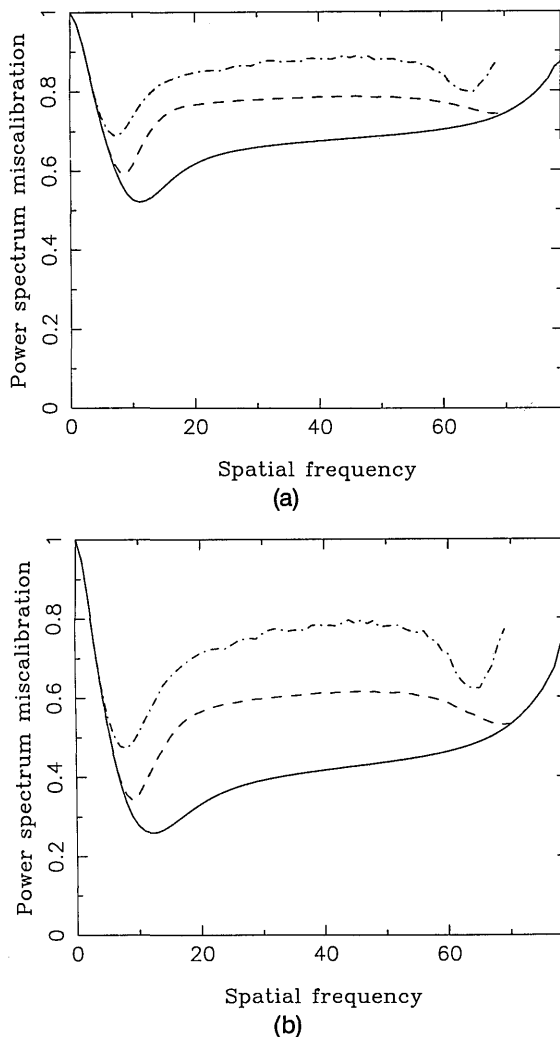
Fig. 5. Fractional miscalibration of the mean power spectrum when the Fried parameter r_0 increases by (a) 20% and (b) 50% between the observations of the source and the calibrator. The solid curves show the miscalibration for the unapodized pupil, the dashed curves show that for the r_0-thick annulus, and the dotted–dashed curves show that for the 0.5-r_0-thick annulus.

clear aperture area, 2.5, 1.6, and 0.75 rad, respectively, confirming that apodization can lower the atmospheric noise on the bispectrum measurements and thus considerably reduce the time taken to image a source. These data suggest that savings in observing time of either ~2.4 or ~11 can be achieved if the two annular pupils are preferred over the unapodized pupil.

3. Amplitude Calibration

An equally important consideration is the sensitivity of the PRM method to the seeing variations. To assess this we calculated the power-spectrum-calibration errors that would occur if r_0 had increased by 20% and 50% between the observations of the source of interest and those of its unresolved calibration star. These results are displayed in Fig. 5, which shows the fractional power spectrum error as a function of spatial frequency. At the lowest spatial frequencies the seeing variations introduce a precipitous drop in calibrated power spectrum amplitude that may incorrectly be interpreted as evidence for an extended component in the observed object, while at higher spatial

frequencies the power spectrum miscalibration is a relatively smooth function of spatial frequency. For both situations considered, the calibration error can be reduced by a factor of 2 if a partially redundant mask is employed. This is particularly noticeable in the low-frequency region of the power spectrum, where the large miscalibration feature caused by the change in the width of the seeing spike is much reduced in height and width. Furthermore, the calibration error is nearly constant over quite a large range of spatial frequencies in the case in which the pupil is apodized, so that the relative amplitudes of most of the Fourier components in the reconstructed image will not be too greatly disturbed by the seeing fluctuations. In contrast, if a fully filled aperture is used these results indicate that there will be a tendency for the calibrated visibility amplitudes to show a gradual rising trend with increasing spatial frequency, which may not be easily recognized as being an artifact of miscalibration.

4. Limiting Magnitudes

Despite the signal-to-noise and calibration advantages described above, the most important drawback of the PRM technique is likely to be its poorer limiting sensitivity when it is compared with filled-aperture speckle methods. To investigate this issue we followed two parallel approaches. The first of these was to evaluate the mean power-spectrum signal level for a source of fixed intensity observed through our three pupil geometries. Because the noise in the power spectrum will be dominated at the lowest light levels by the fixed contribution arising from the readout noise of the detector, the relative magnitudes of the detected power spectrum signal give a direct measure of the relative limiting sensitivity of each pupil configuration. We used the integral derived by Korff[24] to compute the mean power-spectrum-signal level for each of our pupil configurations and plotted the results as a function of spatial frequency in Fig. 6. As expected, the fully filled pupil is superior to the two partially redundant geometries we have examined and provides signal levels
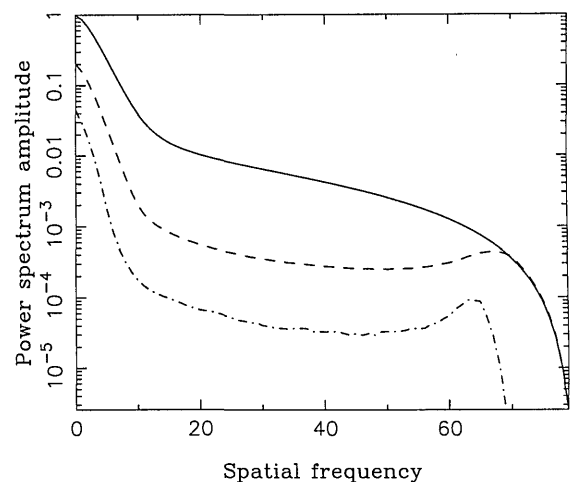

Fig. 6. Mean amplitude of the power spectrum of the specklegrams for the three different pupil geometries discussed in the test when the same unresolved source is observed. The power spectra are shown on the same (arbitrary) scale: The solid curve denotes the unapodized pupil, the dashed curve denotes the r_0-thick annulus, and the dotted–dashed curve denotes the 0.5-r_0-thick annulus.
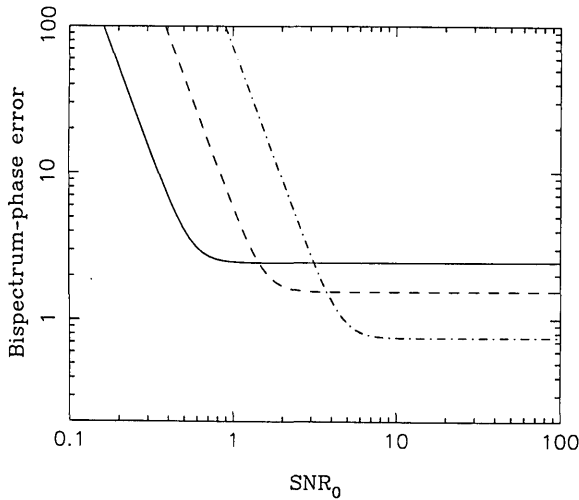
Fig. 7. Median bispectrum-phase error for the observation of a point source as a function of light level. The parameterization of the source flux (SNR$_0$) is described in the text. The solid curve denotes the unapodized pupil, the dashed curve denotes the r_0-thick annulus, and the dotted–dashed curve denotes the 0.5-r_0-thick annulus.

some 16 times higher than those attainable with the annular pupil of width r_0 in the midfrequency region. Remembering that the signal level in the power spectrum is proportional to the square of the flux of the source and assuming a fixed noise level for faint sources, we can translate this into a loss in sensitivity of ~1.5 magnitudes for the r_0-wide, partially redundant pupil when it is compared with the fully filled aperture.

An alternative estimate of the loss in sensitivity was obtained by calculating the dependence of the magnitude of the bispectrum-phase errors as a function of source brightness. This is likely to give a more realistic estimate of the loss in sensitivity when an apodizing mask is used because the bispectrum-phase errors rise much faster than the power spectrum errors as the apparent brightness of the source is reduced. In Appendix A, we calculate the bispectrum-phase error for readout-noise-limited measurements and show it to be given by

$$
\begin{aligned}
\mathrm{var}_\phi[I^{(3)\prime}(\mathbf{u}_1,\mathbf{u}_2)] = {} & \mathrm{var}_\phi[I^{(3)}(\mathbf{u}_1,\mathbf{u}_2)] + \frac{1}{2|\langle I^{(3)}(\mathbf{u}_1,\mathbf{u}_2)\rangle|^2} \\
& \times \{45 N_{\mathrm{pix}}{}^3\sigma^6 + 3 N_{\mathrm{pix}}{}^2\sigma^4[\langle|I(\mathbf{u}_1)|^2\rangle \\
& + \langle|I(\mathbf{u}_2)|^2\rangle + \langle|I(-\mathbf{u}_1-\mathbf{u}_2)|^2\rangle] \\
& + N_{\mathrm{pix}}\sigma^2[\langle|I(\mathbf{u}_1)I(\mathbf{u}_2)|^2\rangle \\
& + \langle|I(\mathbf{u}_2)I(-\mathbf{u}_1-\mathbf{u}_2)|^2\rangle \\
& + \langle|I(-\mathbf{u}_1-\mathbf{u}_2)I(\mathbf{u}_1)|^2\rangle]\},
\end{aligned} \qquad (2)
$$

where $\mathrm{var}_\phi[Q]$ denotes the square of the phase error of the quantity Q, $I^{(3)}(\mathbf{u}_1,\mathbf{u}_2) = I(\mathbf{u}_1)I(\mathbf{u}_2)I(-\mathbf{u}_1-\mathbf{u}_2)$ is the true bispectrum at spatial-frequency coordinates $\mathbf{u}_1$ and $\mathbf{u}_2$, $I^{(3)\prime}(\mathbf{u}_1,\mathbf{u}_2)$ is its measured noisy value, $I(\mathbf{u})$ is the Fourier component at spatial frequency $\mathbf{u}$ of the illumination pattern on the detector, N_{pix} is the number of detector pixels, and σ^2 is the variance of the readout noise per pixel. Figure 7 shows the magnitude of the phase error plotted as a function of light level for the three different pupil geometries for the case of an unresolved source. The curves were calculated by using Eq. (2) together with the

median values of the image-dependent quantities, such as $|\langle I^{(3)\prime}(\mathbf{u}_1,\mathbf{u}_2)\rangle|$, derived from the simulations described in Subsection 2.B.1. The abscissa, which measures the flux detected from the source, was parameterized in terms of a quantity SNR$_0$ given by

$$
\mathrm{SNR}_0 = I_0/\sigma\sqrt{N_{\mathrm{pix}}}, \qquad (3)
$$

where I_0 is the flux detected in one exposure time through an aperture of diameter r_0 and σ and N_{pix} are defined as for Eq. (2). This permits the results of Fig. 7 to be applied to detectors of arbitrary performance. For example, a typical speckle camera with 64 × 64 pixels and a read-out noise of $450e^-$ rms per pixel gives a unit SNR$_0$ for a source of K magnitude of 4.8 if typical observing conditions are assumed ($r_0 = 60$ cm, $\Delta\lambda = 0.20$ μm, $T_{\mathrm{exp}} = 50$ ms). Referring to Fig. 7, we can see that, at high light levels, the phase error is independent of SNR$_0$ and is equal to the value set by the atmospheric noise on the bispectrum, while at low light levels it increases as the inverse cube of the source brightness. This rapid escalation sets a strong limit on the brightness of the faintest source that can be successfully imaged with any of the pupil geometries, a fact that is best appreciated by recognizing that, once below the knee of the SNR curves, a deterioration in the mean seeing parameters of only 10% demands an increase in observing time of approximately a factor of 2. For a given maximum permissible bispectrum-phase error, the low-light-level asymptotes of the curves can be used to derive losses of one and two magnitudes in limiting sensitivity for the r_0 and $r_0/2$-thickness annuli when compared with the unapodized pupil.

In summary, we conclude that the use of either of the annular pupils will lead to a loss of perhaps one or two magnitudes of sensitivity at most and so would be suitable for studies of all but the faintest or most-resolved sources.

3. IMAGE RECOVERY

The methods for image recovery from speckle interferograms obtained by using annular partially redundant pupils have been discussed in some detail by Haniff et al.,[15,25] and so for the sake of economy we present only a brief discussion of the most important issues here. In line with most-recent developments, we concentrate attention on the recovery of the source Fourier amplitudes and phases from the power spectrum, $|I(\mathbf{u})|^2$, and the bispectrum, $I^{(3)}(\mathbf{u}_1,\mathbf{u}_2)$, of the image, respectively.

Because the transfer function of any annular pupil samples the whole of the Fourier plane to the limit set by the outer diameter of the annulus, the procedures for visibility amplitude recovery are unchanged in the presence of an annular apodizing mask. On the other hand, in the presence of atmospheric turbulence the support of the bispectral transfer function of an annular pupil does vary with the annular width. Referring to the heuristic model of Subsection 2.A.1, we can see that this peculiar behavior arises because a useful signal can be measured only at bispectrum points that correspond to physically realizable triangles in the telescope pupil. If this is not the case, the component of the bispectrum that can be averaged coherently between exposures, i.e., the part that depends on a phase-closure relation's being satisfied, vanishes.[22,26,27] Of course, the spatial correlations of the real

wave-front perturbations allow additional, nonclosing triplets of baselines to contribute so long as the nonclosing components are $\leq r_0$ in magnitude.

For example, for an annulus of infinitesimal width most of the bispectrum points contain no useful information at all, even when the individual Fourier components can be measured with an infinite SNR. For the annuli considered here, the reduction of the bispectral support is relatively unimportant because their widths are still quite large; the volumes of the nondegenerate portion of the bispectrum sampled by the unapodized aperture and the annuli of widths r_0 and $r_0/2$ are in the ratio 1:0.78:0.29. However, it is important that such nonphysical bispectral measurements not be used during the Fourier-phase-retrieval procedure. To date two indexing algorithms for the generation of nondegenerate subsets of bispectrum coordinates have been presented in the literature.[28,29] Neither of these schemes explicitly includes tests to check whether a bispectral triangle can be realized in the telescope pupil, but it is straightforward to include such tests that use either formalism, to ensure that valuable resources are not wasted during the accumulation of the bispectrum measurements.

Fourier-phase retrieval can then be accomplished by using either recursive[30] or global least-squares[29,31-33] algorithms. In the former case, care must be taken to construct a continuous recursive path from the low-frequency starting points to all higher spatial frequencies; simple-minded schemes, e.g., straight-line paths from the origin, may require bispectral components that cannot be measured. Finally, the image can be reconstructed from the recovered Fourier amplitudes and phases, together with any other *a priori* constraints, such as positivity, by using standard methods.

4. IMAGING SENSITIVITY AND FIDELITY

Because of the relative immaturity of phase-retrieval methods in speckle interferometry, most previous high-resolution, near-infrared studies have relied on model-fitting techniques to establish the properties of the source under study from visibility amplitude data alone. We have discussed the errors that may occur in these data, but, in view of the current concern with true imaging from speckle data, we thought it vital to investigate the propagation of these errors in the reconstructed images. Images are the most natural and immediate way in which the information in the Fourier data can be conveyed, but the consequences of measurement errors are more complex in the image plane than in the aperture plane. An intuitive appreciation of these consequences is perhaps best gained from the examination of sample reconstructions.

We have already seen that the visibility amplitude data are liable to large-yet-subtle errors if the local seeing conditions vary on moderate time scales. As is evident from Fig. 5, these will often lead to the spurious detection of large-scale features that are apparently resolved out by the interferometer on relatively short baselines. This type of reconstruction error, introduced by a systematic misfit between the measured and the true Fourier data, is potentially far more hazardous than any increase in the level of random noise associated with the visibility measurements. Under these circumstances the use of a quan-

tity such as the dynamic range in the recovered image, i.e., the ratio of the peak flux to the rms noise level in empty regions of the map, is no longer adequate to quantify the success of the image recovery. A particular point of concern is the situation in which the recovered image has low rms fluctuations in the source-free parts of the field but does not reproduce the true source structure accurately. In the parlance of radio astronomy such images have high dynamic range but low fidelity.

To investigate the effects of pupil redundancy on the dynamic range and the fidelity of the interferometric imaging procedure, we have once again made use of numerical simulations. Using the methods outlined in Section 2, we generated simulated specklegrams and attempted to reconstruct images by using the accumulated power spectrum and bispectrum. As a test source we chose a well-resolved object with dimensions of $\sim 5 \times 12$ beam widths, which we believe is sufficiently complex to provide a realistic simulation of an astrophysically interesting target. A contour plot of the model source is shown in Fig. 8. Sequences of 512 point-source and test-source specklegrams were generated for each of the three pupil configurations, assuming a standard value of D/r_0 of 8, and for two different light levels. These corresponded to SNR_0 values of ∞ and 3.7, respectively. Point-source speckle interferograms were also generated by using atmospheric phase screens appropriate to seeing conditions 20% and 50% better than those used to generate the test source data so as to mimic the effects of 20% and 50% seeing miscalibrations.

For the simulations presented here we have taken the simplest possible approach to visibility-amplitude recovery and obtained the calibrated visibility amplitudes by dividing the normalized source power spectrum by that of the calibration star and taking the square root of the resulting quantity.

Rather than recovering the source Fourier phases from all the measured bispectrum points by using a least-
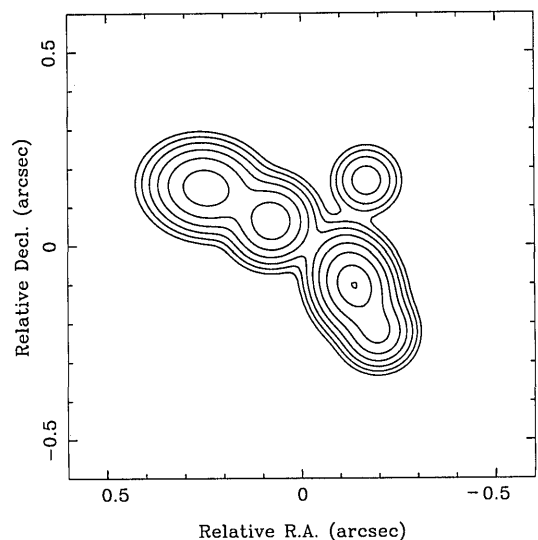

Fig. 8. Contour plot of the test source used in the imaging simulations. The contour levels are at 1, 2, 4, 8, 16, 32, 64, and 99% of the peak flux. The object has been convolved with a Gaussian beam with a FWHM of 0.083 arcsec to facilitate comparison with the image reconstructions.

squares method, we have found it advantageous to use a simpler technique that employs only a small subset of the bispectral data. This method is described in detail by Haniff et al.,[15,25] and so only a brief résumé is presented here. The procedure involves sampling the Fourier plane at a discrete set of locations that correspond to the Fourier coverage of an interferometric array, or pseudo-array. If the pseudoarray elements are chosen so as to lie within the clear area of the telescope pupil, the bispectrum points selected are guaranteed to correspond to physically realizable triangles. So long as the Fourier plane is adequately sampled, the exact location and number of the pseudoarray elements are relatively unimportant. For all these simulations we used an array comprising 17 elements, thereby selecting 136 Fourier and 680 bispectrum components. This method of sampling the Fourier plane permits the direct use of the radio astronomical self-calibration technique[34] to recover images of the source. This has the advantage that a priori image-plane constraints (source positivity and finite extent) are incorporated explicitly during the Fourier-phase retrieval. These constraints also provide a measure of resistance to miscalibration by allowing the a priori information to counterbalance the systematic errors in the visibility data. The effects of variations in the seeing on the reconstructed images are likely to be further reduced because the sparse sampling of the bispectral data implicitly rejects the most poorly calibrated Fourier data, i.e., the lowest-spatial-frequency components (see Subsection 2.A.3).

Thus by using a method that attempts to minimize the miscalibration-induced artifacts in our images we can hope to estimate conservatively when the fidelity of the image deteriorates to an unacceptable level. An additional feature of our method is that its computational requirements are significantly fewer than those of methods that utilize the full bispectrum. This is especially important here because of the large number of image reconstructions we have performed.

At low signal levels such a procedure will, however, become less than optimal. At low light levels the locations of useful signals in the bispectrum contract toward the near-axis region, i.e., the region corresponding to baseline triangles that include at least one side of length $\lesssim r_0/\lambda$. With a self-calibration-based scheme it is difficult to sample this region preferentially, and so the majority of the accumulated bispectrum data are likely to have low SNR's. A more-complete discussion of the rationale for using self-calibration-based methods and their relationship to more-standard procedures for image recovery from speckle data can be found in Ref. 25 and the references therein. We reiterate that the selection of this particular image-reconstruction algorithm was made in the interests of expediency, and so our choice should not be taken to represent an optimum strategy.

The results of our simulations are shown in Figs. 9 and 10. In each figure, the sequence from left to right corresponds to decreasing pupil redundancy and the sequence from top to bottom corresponds to increasing seeing miscalibration. At the high-light-level limit and under conditions of good seeing stability [Figs. 9(a)–9(c)], all the recovered images are of high dynamic range and fidelity, showing noise features at levels of order 1% of their peak intensity. There is marginal evidence that the pupil ge-

ometry of lowest redundancy provides the best map, but there is little difference between the fully filled aperture image and its r_0-wide-annulus counterpart. This similarity in map quality despite the large differences between the median high-light-level bispectrum-phase errors of the different pupil geometries (see Subsection 2.B.2) reflects the fact that the properties of the reconstructed image are dominated by atmospheric noise on the amplitudes.

The differences between the pupil geometries become more readily apparent as the seeing stability deteriorates. With a 20% miscalibration the images show a monotonic enhancement in fidelity and noise suppression as the pupil redundancy decreases, in accord with the predictions of Subsection 2.A. The major qualitative result of this level of miscalibration is simply to amplify the number of spurious features over the field of view, although the image of the fully filled aperture [Fig. 9(d)] is beginning to show signs of a weak halo surrounding the test object.

Finally, at the highest levels of miscalibration, the more-insidious consequences of seeing variations become clear: With the unapodized pupil [Fig. 9(g)] the whole of the image sits on a low-level plateau of emission at approximately 1% of the peak flux and the major noise components in the map cluster around the central object. In addition, the edges of the test source appear to be less steep than they really are, so that the overall outcome is to smooth the image and raise it on a shallow pedestal of emission. With the r_0-wide annular pupil [Fig. 9(h)] these effects are much reduced, although the noise features seen at the 2% level still sit on shallow ridges of lower-level emission that are spurious. Only the data obtained through the least-redundant pupil are sufficiently unaffected to permit faithful image recovery. Despite the 50% calibration errors, the reconstructed image is well shaped, has high contrast, and correctly reproduces all the features in the true object. The main artifacts in the image are the small number of spiky noise peaks that appear randomly across the field, features that are easily recognized as being unreal.

For lower light levels our conclusions are somewhat changed. As is clear from the right-hand column of Fig. 10, none of the inversions from data obtained with the $r_0/2$-sized pupil was successful. The detected flux level in this set of simulations is too low to permit a reliable estimation (i.e., with a SNR of more than 2) of the source bispectrum phases to be made for more than 40% of the subset of bispectrum points used by our mapping procedure. It is a well-established result in radio astronomy that the success of self-calibration methods is contingent on most of the bispectrum phases' having SNR's of $\gtrsim 2$, and so this is an understandable failure. The other two pupil geometries, having larger effective areas, do not suffer from this problem but serve to confirm the results shown in Fig. 9. In the absence of any seeing miscalibration the filled aperture and successful PRM map [Figs. 10(a) and 10(b)] are quite similar, there being tentative evidence in the shape of the faintest image component to suggest that the latter is slightly more reliable than the former. An important point is that although the rms noise levels in the regions away from the central source are quite low, i.e., well below the 2% level of the lowest contour plotted in the figure, the maps are not reliable at

that level because the locations of the outermost contours are not in agreement with those of the model (cf. Fig. 8). Thus even at this source brightness the fidelity of the maps is lower than the dynamic range suggests.

With increasing miscalibration the filled pupil map becomes progressively noisier and less reliable, until finally under the most unstable seeing conditions investigated [Fig. 10(g)] the mapping procedure fails. This does not occur for the data recorded through the r_0-wide, partially redundant pupil, which can still be inverted and provides a useful, although noisy, image reconstruction. This fail-

ure of the fully filled aperture inversion can be understood in terms of a combination of the effects of noise on the raw Fourier data, together with the incompatibility between the miscalibrated visibility amplitudes, the less-affected bispectrum phases, and the assumptions inherent in the CLEAN algorithm utilized in our self-calibration loop. We remind the reader that the mean values of the measured bispectrum phases are not affected in any systematic fashion by the seeing fluctuations. These results confirm that, in this instance, the use of a partially redundant pupil provides the key to maintaining a useful
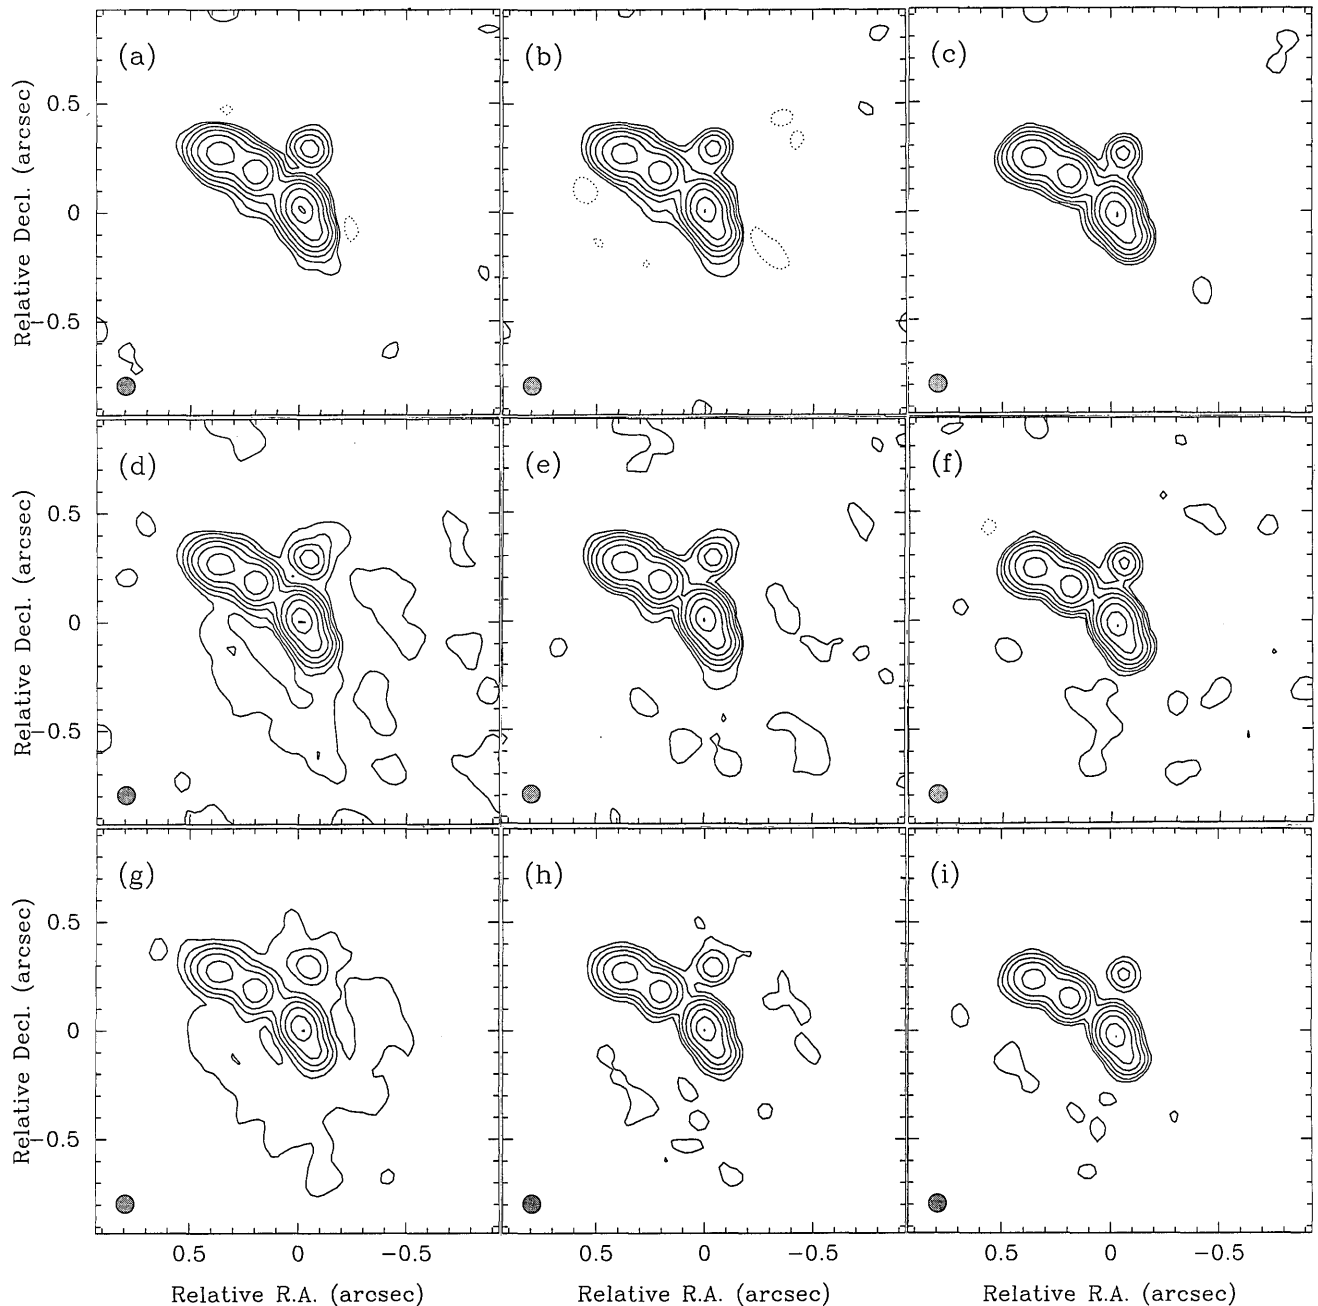

Fig. 9. Images reconstructed from simulated specklegrams at infinite light level. A self-calibration-based method was used to recover the images from a total of 136 Fourier amplitudes and 680 bispectrum points. The maps are arranged from left to right in order of decreasing pupil area, i.e., the unapodized aperture, the r_0-thick annulus, and the 0.5-r_0-thick annulus, respectively, and from top to bottom in order of increasing seeing variation between source and calibrator data sets, namely 0, 20, and 50% variation, respectively. The contour levels plotted are at -4, -2, -1, 1, 2, 4, 8, 16, 32, 64, and 99% of the peak flux, except in the lower three maps, for which the -1 and 1% contours have been removed. Negative contour levels are shown dashed. The CLEAN beam used (FWHM = 0.083 arcsec) is shown in the bottom left-hand corner of each image.
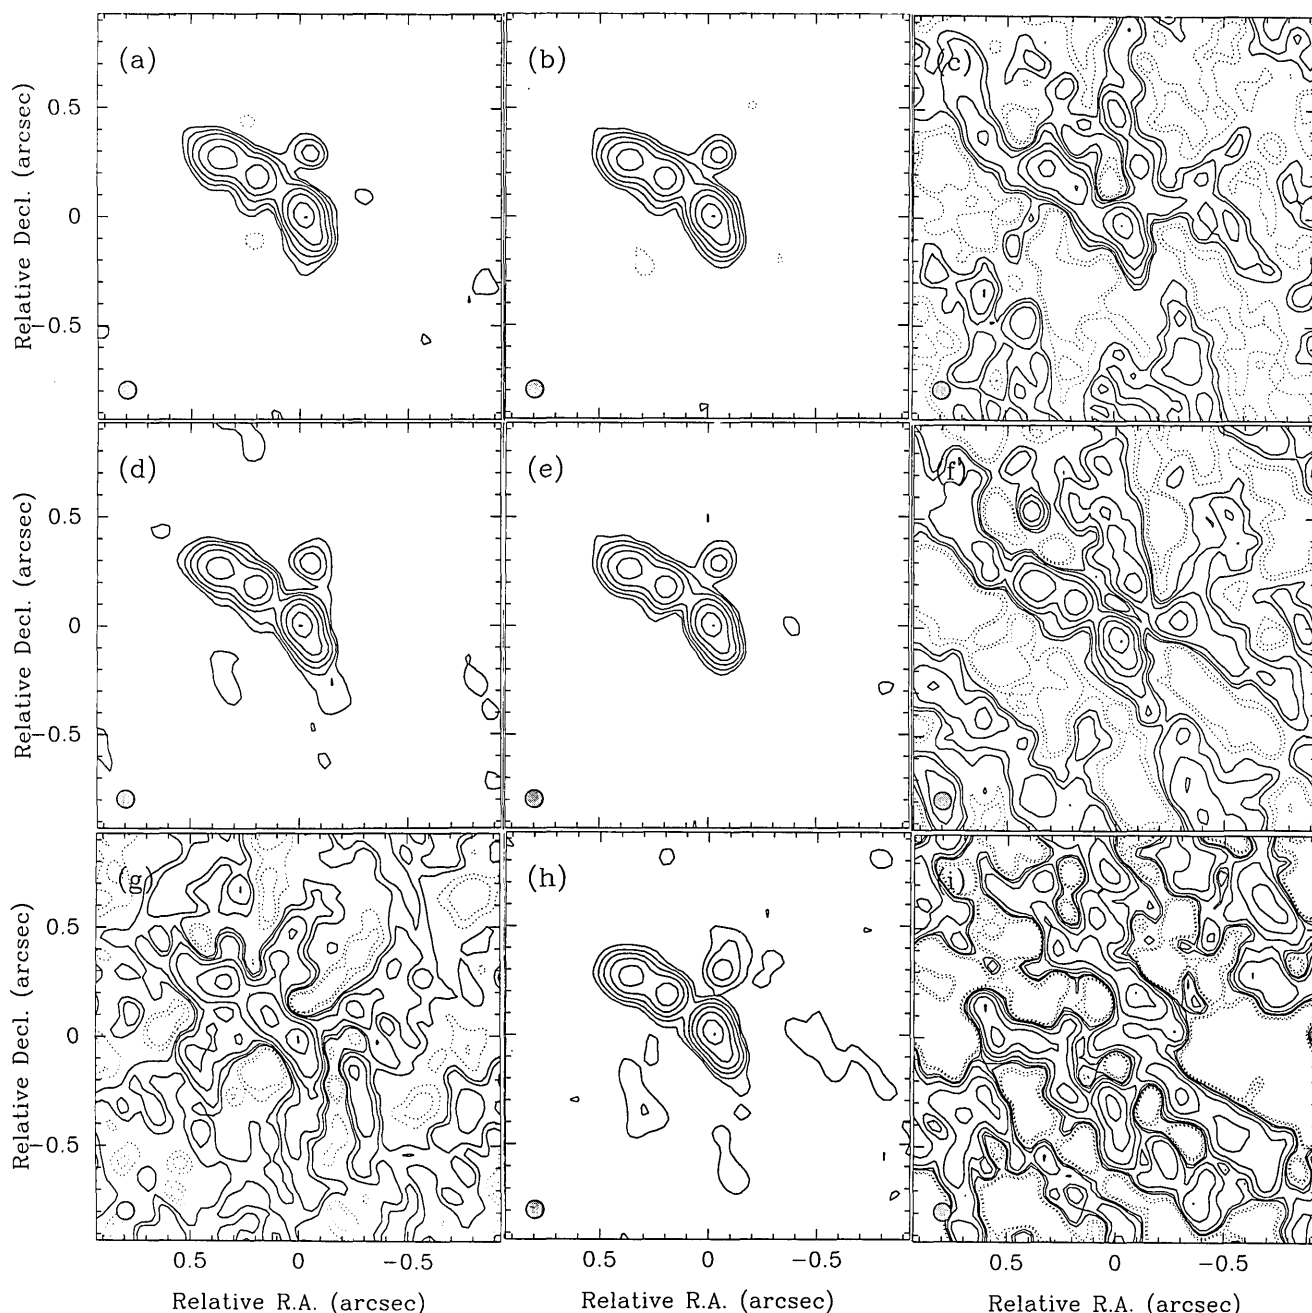
Fig. 10.   Images reconstructed from simulated specklegrams at a low light level, corresponding to a SNR_0 of 3.7.   The same method and subsets of visibility amplitude and bispectral data as in Fig. 9 were used, and the maps are arranged in the same order with regard to pupil redundancy and seeing miscalibration.   The contour levels plotted are at -4, -2, 2, 4, 8, 16, 32, 64, and 99% of the peak flux, with negative contours shown dashed.

imaging capability in the presence of the seeing variations that are likely to be experienced to some degree at all observing sites.

We must stress that the image reconstructions presented here should be interpreted with caution. Because the Fourier-plane sampling of our raw data is not complete and samples the lowest spatial frequencies only sparsely, the propagation of the dominant effects of the seeing miscalibration, which occur in the seeing spike at the lowest spatial frequencies, is suppressed before it can influence the reconstructions.   On the other hand, the sparsity of the Fourier-plane coverage does mean that a sizable fraction of the data in the raw specklegrams must be discarded.   Fourier-inversion schemes that utilize all the

available Fourier data are likely to be more successful in restoring images from noisy data than the sparsely sampled method we have employed here.   Indeed, we have verified that the poorly calibrated Fourier data sets of Figs. 10(g)–10(i) can *all* be inverted with a bispectrum code that utilizes all the measured power spectrum data and only a moderate fraction (~10%) of the available bispectral data.   These alternative image reconstructions, which are displayed in Fig. 11, while appearing successful, still serve to confirm the results presented above.   At the highest level of pupil redundancy the recovered image appears to be of extremely high quality and shows no obvious noise features above a level of 0.5% of the peak flux.   However, the restored image sits on a broad plateau of
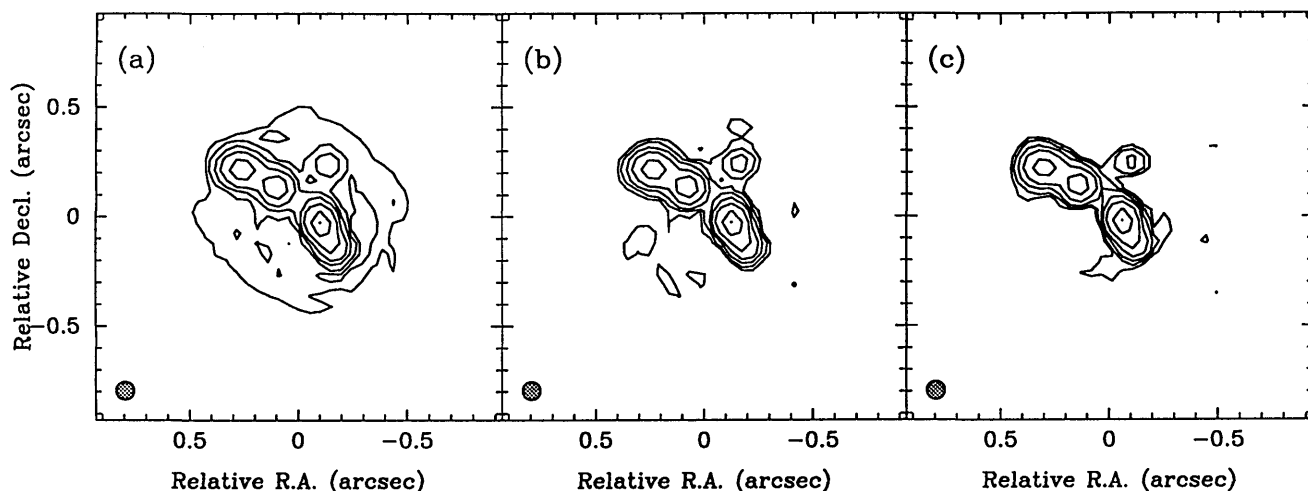
Fig. 11. Images recovered from the same sequences of specklegrams as those used to prepare Figs. 10(g)–10(i) but by use of a least-squares bispectrum algorithm. The reconstructions utilized 1060 Fourier amplitudes and 21,685 bispectrum points. As before, the maps are presented in order of decreasing pupil area from left to right. The contour levels are identical to those in Fig. 10.

emission that has a maximum value of approximately 5% of the peak flux that, while appearing to be physically plausible, is nothing other than an artifact of the systematic visibility-amplitude errors. As the pupil redundancy is reduced the integrated flux in the plateau falls and the fidelity of the recovered image increases, until finally in the case of the data obtained through the $r_0/2$-wide annulus the only vestiges of the miscalibration are a number of diffuse patches of weak emission that are at the same level (~1%) as the noise peaks introduced because of the lower SNR of the data. We conclude that, despite the improved resistance to *noise* that this alternative mapping procedure affords, our main results concerning the *fidelity* of the imaging procedure are independent of the particular image-reconstruction algorithm adopted.

5. SUMMARY

In this paper we have examined the effects of pupil redundancy on the reliability and the sensitivity of speckle imaging at near-infrared wavelengths. We have argued that careful consideration should go into the choice of the pupil geometry since it can have a significant effect on the resulting data. For one particular geometry, namely, an annular mask, we find that considerable improvements on the default fully filled geometry can be achieved: The atmospheric noise on the bispectrum phases can be substantially reduced, as can the magnitude of the perturbations to the visibility-amplitude calibration caused by various seeing conditions. The latter property is particularly noteworthy because the visibility data can be corrupted in a manner that is a smooth function of spatial frequency and this corruption may therefore not be easily identified. This can lead to misinterpretation of the visibility data and of the images reconstructed from these data. As is shown by our simulations, the dynamic range of the recovered maps may give a misleading impression as to their fidelity when miscalibration effects are important.

The major disadvantage of this geometry is a moderate loss in sensitivity at the lowest light levels, corresponding to a factor of ~4 (1.5 magnitudes). However, in many astronomical implementations of speckle interferometry this consequence is likely to be outweighed by the resistance to seeing miscalibration that such a pupil function affords. For moderately bright near-infrared sources the use of alternative partially redundant pupil geometries should permit substantial improvements in imaging fidelity.

APPENDIX A: NOISE ON THE BISPECTRUM FOR READOUT-NOISE-LIMITED MEASUREMENTS

In this calculation we shall ignore the effects of photon noise, which will be assumed to be negligible in comparison with the noise contribution from the detector itself. A similar calculation was performed by Roddier,[35] who evaluated the SNR for bispectral measurements corrupted by additive noise. Our treatment differs from his in two ways: We characterize the reliability of a bispectrum measurement in terms of the phase error rather than of the SNR and we explicitly include the effects of correlations between the noise in the measured Fourier components that result from a single detector being used to measure all the Fourier components of the image.

We consider an array detector with zero mean-additive-readout noise:

$$i'(\mathbf{x}_p) = i(\mathbf{x}_p) + n(\mathbf{x}_p),$$

where $i(\mathbf{x}_p)$ is the flux incident upon pixel p at angular coordinate $\mathbf{x}_p$ in the focal plane of the telescope, $n(\mathbf{x}_p)$ is the readout noise for that pixel, and $i'(\mathbf{x}_p)$ is the measured noisy signal. We assume that the noise upon any pixel is independent of both the flux incident upon that pixel and the noise on any other pixel, and for simplicity we further assume that the noise has a normal probability distribution. The latter proviso is not crucial to our argument, and analogous results can be derived straightforwardly for other symmetric distributions. These constraints lead to the following results, which are used below:

$$\langle i(\mathbf{x}_p)n(\mathbf{x}_q)\rangle = 0 \qquad \forall p, q, \tag{A1}$$

$$\langle n(\mathbf{x}_p)n(\mathbf{x}_q)\rangle = \delta_{pq}\sigma^2, \tag{A2}$$

$$\langle n(\mathbf{x}_p)^{2k+1}\rangle = 0 \qquad k = 0, 1, 2, \ldots, \tag{A3}$$

where δ_{pq} is the Kronecker delta and σ^2 is the variance of the noise, assumed to be equal for all the pixels.

We now define the bispectrum measurement:

$$I^{(3)\prime}(\mathbf{u}_1, \mathbf{u}_2) \equiv I'(\mathbf{u}_1)I'(\mathbf{u}_2)I'(-\mathbf{u}_1 - \mathbf{u}_2),$$

where $I'(\mathbf{u}) = \Sigma_p \exp(2\pi i\mathbf{u} \cdot \mathbf{x}_p)i'(\mathbf{x}_p)$ represents the Fourier component at spatial frequency $\mathbf{u}$ of the detected specklegram. First let us evaluate the mean value of this quantity to show that it is an unbiased estimate of the true bispectrum and also to illustrate the steps used in evaluating the variance:

$$\langle I^{(3)\prime}(\mathbf{u}_1, \mathbf{u}_2)\rangle = \sum_{p,q,r} \exp\{2\pi i[\mathbf{u}_1 \cdot \mathbf{x}_p + \mathbf{u}_2 \cdot \mathbf{x}_q$$
$$- (\mathbf{u}_1 + \mathbf{u}_2) \cdot \mathbf{x}_r]\} \langle [i(\mathbf{x}_p) + n(\mathbf{x}_p)]$$
$$\times [i(\mathbf{x}_q) + n(\mathbf{x}_q)][i(\mathbf{x}_r) + n(\mathbf{x}_r)]\rangle$$

$$= \sum_{p,q,r} \exp\{2\pi i[\mathbf{u}_1 \cdot \mathbf{x}_p + \mathbf{u}_2 \cdot \mathbf{x}_q$$
$$- (\mathbf{u}_1 + \mathbf{u}_2) \cdot \mathbf{x}_r]\} [\langle i(\mathbf{x}_p)i(\mathbf{x}_q)i(\mathbf{x}_r)\rangle$$
$$+ \langle i(\mathbf{x}_p)i(\mathbf{x}_q)n(\mathbf{x}_r)\rangle + \langle i(\mathbf{x}_q)i(\mathbf{x}_r)n(\mathbf{x}_p)\rangle$$
$$+ \langle i(\mathbf{x}_r)i(\mathbf{x}_p)n(\mathbf{x}_q)\rangle + \langle i(\mathbf{x}_p)n(\mathbf{x}_q)n(\mathbf{x}_r)\rangle$$
$$+ \langle i(\mathbf{x}_q)n(\mathbf{x}_r)n(\mathbf{x}_p)\rangle + \langle i(\mathbf{x}_r)n(\mathbf{x}_p)n(\mathbf{x}_q)\rangle$$
$$+ \langle n(\mathbf{x}_p)n(\mathbf{x}_q)n(\mathbf{x}_r)\rangle].$$

Making use of Eqs. (A1)–(A3), we find that

$$\langle I^{(3)\prime}(\mathbf{u}_1, \mathbf{u}_2)\rangle = \langle I(\mathbf{u}_1)I(\mathbf{u}_2)I(-\mathbf{u}_1 - \mathbf{u}_2)\rangle + \langle I(\mathbf{u}_1)\rangle$$
$$\times \sum_r \sigma^2 \exp(-2\pi i\mathbf{u}_1 \cdot \mathbf{x}_r) + \langle I(\mathbf{u}_2)\rangle$$
$$\times \sum_p \sigma^2 \exp(-2\pi i\mathbf{u}_2 \cdot \mathbf{x}_p) + \langle I(-\mathbf{u}_1 - \mathbf{u}_2)\rangle$$
$$\times \sum_q \sigma^2 \exp[2\pi i(\mathbf{u}_1 + \mathbf{u}_2) \cdot \mathbf{x}_q],$$

where $I(\mathbf{u}) = \Sigma_p \exp(2\pi i\mathbf{u} \cdot \mathbf{x}_p)i(\mathbf{x}_p)$ is the Fourier component at spatial frequency $\mathbf{u}$ of the true (i.e., noise-free) specklegram. Clearly the first term in the above equation is the bispectrum of the true specklegram $I^{(3)}(\mathbf{u}_1, \mathbf{u}_2)$. The other terms will be zero if at least one of two conditions is met: First, if the detector array is a regular rectangular grid and the spatial frequencies $\mathbf{u}_1$ and $\mathbf{u}_2$ are chosen to give an integral number of cycles across the detector, then the summations in the above equation will be zero. Note that this depends on the assumption that the noise is uniform from pixel to pixel: In real detectors the presence of hot pixels and other nonuniformities will mean that the summations have a small but finite value. Alternatively these terms will be zero if the mean values of the specklegram Fourier components $\langle I(\mathbf{u})\rangle$ are zero. This will be true if their phases vary randomly between exposures, a condition that is satisfied for all but low spatial frequencies where $|\mathbf{u}| < r_0/\lambda$. The combination of these two conditions means that in practice there will be a small amount of bias for bispectrum points where one or more of the component spatial frequencies have a small magnitude. This bias is likely to be important only at very low SNR's, so we shall ignore it here and write simply

$$\langle I^{(3)\prime}(\mathbf{u}_1, \mathbf{u}_2)\rangle \approx \langle I^{(3)}(\mathbf{u}_1, \mathbf{u}_2)\rangle. \qquad (A4)$$

We now come to the evaluation of the noise on the bispectrum measurement. Because the bispectrum is a complex quantity, its variance will be a function of direction in the complex plane. A useful formalism for quantifying such variances has been presented by Buscher,[36] who shows that the variance of a complex quantity Q in any direction in the complex plane can be derived from two parameters, $\text{var}_1(Q)$ and $\text{var}_2(Q)$, which have the following definitions:

$$\text{var}_1(Q) \equiv \langle|Q|^2\rangle - |\langle Q\rangle|^2, \qquad (A5)$$

$$\text{var}_2(Q) \equiv \langle(Q)^2\rangle - (\langle Q\rangle)^2, \qquad (A6)$$

so that

$$\text{var}(Q, \theta) = \tfrac{1}{2}\{\text{var}_1(Q) + \Re[\exp(-i2\theta)\text{var}_2(Q)]\}, \qquad (A7)$$

where $\text{var}(Q, \theta)$ is the variance of Q along a direction at an angle θ to the real axis and $\Re(x)$ denotes the real part of the complex quantity x.

We consider initially the evaluation of the first term of $\text{var}_1[I^{(3)\prime}(\mathbf{u}_1, \mathbf{u}_2)]$:

$$\langle|I^{(3)\prime}(\mathbf{u}_1, \mathbf{u}_2)|^2\rangle = \sum_{k,l,m,p,q,r} \exp\{2\pi i[\mathbf{u}_1 \cdot \mathbf{x}_k + \mathbf{u}_2 \cdot \mathbf{x}_l$$
$$- (\mathbf{u}_1 + \mathbf{u}_2) \cdot \mathbf{x}_m - \mathbf{u}_1 \cdot \mathbf{x}_p - \mathbf{u}_2 \cdot \mathbf{x}_q$$
$$+ (\mathbf{u}_1 + \mathbf{u}_2) \cdot \mathbf{x}_r]\} \langle [i(\mathbf{x}_k) + n(\mathbf{x}_k)]$$
$$\times [i(\mathbf{x}_l) + n(\mathbf{x}_l)][i(\mathbf{x}_m) + n(\mathbf{x}_m)]$$
$$\times [i(\mathbf{x}_p) + n(\mathbf{x}_p)][i(\mathbf{x}_q) + n(\mathbf{x}_q)]$$
$$\times [i(\mathbf{x}_r) + n(\mathbf{x}_r)]\rangle. \qquad (A8)$$

As when calculating the mean of the bispectrum, we expand the bracketed factors in Eq. (A8) and substitute from Eqs. (A1)–(A3). A large fraction of the terms in this expansion can be eliminated if either

$$\sum_p \exp(2\pi i\mathbf{u} \cdot \mathbf{x}_p) = 0$$

$$\forall \mathbf{u} : \pm\mathbf{u} \in \{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_1 + \mathbf{u}_2, \mathbf{u}_1 - \mathbf{u}_2, 2\mathbf{u}_1 + \mathbf{u}_2, \mathbf{u}_1 + 2\mathbf{u}_2\} \qquad (A9)$$

or

$$\langle I(\mathbf{u}_a)I(\mathbf{u}_b)\ldots I(\mathbf{u}_c)\rangle = 0$$

$$\forall \mathbf{u}_a, \mathbf{u}_b, \ldots, \mathbf{u}_c : \mathbf{u}_a + \mathbf{u}_b + \ldots + \mathbf{u}_c \neq 0. \qquad (A10)$$

These conditions will be satisfied under similar circumstances as before; i.e., either the detector is uniform or $|\mathbf{u}| \gg r_0/\lambda$ $\forall \mathbf{u} : \pm\mathbf{u} \in \{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_1 + \mathbf{u}_2, \mathbf{u}_1 - \mathbf{u}_2, 2\mathbf{u}_1 + \mathbf{u}_2, \mathbf{u}_1 + 2\mathbf{u}_2\}$. If in addition the number of pixels in the detector array is large ($N_{\text{pix}} \gg 1$), Eq. (A8) then reduces to

$$\langle|I^{(3)\prime}(\mathbf{u}_1, \mathbf{u}_2)|^2\rangle \approx 45N_{\text{pix}}^3\sigma^6 + 3N_{\text{pix}}^2\sigma^4[\langle|I(\mathbf{u}_1)|^2\rangle$$
$$+ \langle|I(\mathbf{u}_2)|^2\rangle + \langle|I(-\mathbf{u}_1 - \mathbf{u}_2)|^2\rangle]$$
$$+ N_{\text{pix}}\sigma^2[\langle|I(\mathbf{u}_1)I(\mathbf{u}_2)|^2\rangle$$
$$+ \langle|I(\mathbf{u}_2)I(-\mathbf{u}_1 - \mathbf{u}_2)|^2\rangle$$
$$+ \langle|I(-\mathbf{u}_1 - \mathbf{u}_2)I(\mathbf{u}_1)|^2\rangle]$$
$$+ \langle|I^{(3)\prime}(\mathbf{u}_1, \mathbf{u}_2)|^2\rangle. \qquad (A11)$$

The calculation for $\langle[I^{(3)'}(\mathbf{u}_1, \mathbf{u}_2)]^2\rangle$ proceeds in the same manner and yields

$$\langle[I^{(3)'}(\mathbf{u}_1, \mathbf{u}_2)]^2\rangle \approx \langle[I^{(3)}(\mathbf{u}_1, \mathbf{u}_2)]^2\rangle, \qquad (A12)$$

which we note is independent of the presence of the read-out noise.

Combining relations (A4)–(A6), Eq. (A7), and relations (A11) and (A12), we find that the variance of the bispectrum in an arbitrary direction in the complex plane is given by

$$
\begin{aligned}
\mathrm{var}[I^{(3)'}(\mathbf{u}_1, \mathbf{u}_2), \theta] = {} & \tfrac{1}{2}\{45 N_{\mathrm{pix}}^3 \sigma^6 + 3 N_{\mathrm{pix}}^2 \sigma^4 [\langle|I(\mathbf{u}_1)|^2\rangle \\
& + \langle|I(\mathbf{u}_2)|^2\rangle + \langle|I(-\mathbf{u}_1 - \mathbf{u}_2)|^2\rangle] \\
& + N_{\mathrm{pix}} \sigma^2 [\langle|I(\mathbf{u}_1)I(\mathbf{u}_2)|^2\rangle \\
& + \langle|I(\mathbf{u}_2)I(-\mathbf{u}_1 - \mathbf{u}_2)|^2\rangle \\
& + \langle|I(-\mathbf{u}_1 - \mathbf{u}_2)I(\mathbf{u}_1)|^2\rangle]\} \\
& + \mathrm{var}[I^{(3)}(\mathbf{u}_1, \mathbf{u}_2), \theta]. \qquad (A13)
\end{aligned}
$$

The last term in this expression arises from the fluctuations in the high-light-level image that are due to the effects of atmospheric turbulence, while the remaining terms are a consequence of the detector noise.

The square of the phase error of a quantity Q is computed as

$$\mathrm{var}_\phi(Q) = \frac{\mathrm{var}[Q, \theta = \arg(\langle Q\rangle) + \pi/2]}{|\langle Q\rangle|^2}, \qquad (A14)$$

and so by combining Eqs. (A13) and (A14) we obtain Eq. (2) in the main text.

ACKNOWLEDGMENTS

REFERENCES

1. I. Gatley, D. L. DePoy, and A. M. Fowler, "Astronomical imaging with infrared array detectors," Science **242**, 1264–1270 (1988).
2. See, e.g., *Infrared Astronomy with Arrays*, C. G. Wynn-Williams, E. E. Becklin, and L. H. Good, eds. (Institute for Astronomy, University of Hawaii, Honolulu, Hawaii, 1988).
3. J. M. Beckers, J. C. Christou, R. G. Probst, S. T. Ridgway, and O. von der Lühe, "First results with the NOAO 2-d speckle camera for infrared wavelengths," in *High Resolution Imaging by Interferometry*, F. Merkle, ed., Vol. 29 of European Southern Observatory Conference Proceedings (European Southern Observatory, Garching bei München, Germany, 1988), pp. 393–400.
4. J. C. Christou, S. T. Ridgway, D. F. Buscher, C. A. Haniff, and D. W. McCarthy, Jr., "High spatial resolution imaging of circumstellar envelopes in the near infrared," in *Astrophysics with Infrared Arrays*, R. Elston, ed., Vol. 14 of ASP Conference Series (Astronomical Society of the Pacific, Tucson, Ariz., 1991), pp. 133–138.
5. J. C. Dainty and A. H. Greenaway, "Estimation of spatial power spectra in speckle interferometry," J. Opt. Soc. Am. **69**, 786–790 (1979).
6. T. Nakajima, "Signal-to-noise ratio of the bispectral analysis of speckle interferometry," J. Opt. Soc. Am. A **5**, 1477–1491 (1988).
7. G. R. Ayers, M. J. Northcott, and J. C. Dainty, "Knox–Thompson and triple correlation imaging through atmospheric turbulence," J. Opt. Soc. Am. A **5**, 963–985 (1988).
8. W. T. Rhodes and J. W. Goodman, "Interferometric technique for recording and restoring images degraded by unknown aberrations," J. Opt. Soc. Am. **63**, 647–657 (1973).
9. J. E. Baldwin, C. A. Haniff, C. D. Mackay, and P. J. Warner, "Closure phase in high-resolution optical imaging," Nature (London) **320**, 595–597 (1986).
10. C. A. Haniff, C. D. Mackay, D. J. Titterington, D. Sivia, J. E. Baldwin, and P. J. Warner, "The first images from optical aperture synthesis," Nature (London) **328**, 694–696 (1988).
11. T. Nakajima, S. R. Kulkarni, P. W. Gorham, A. M. Ghez, G. Neugebauer, J. B. Oke, T. A. Prince, and A. C. S. Readhead, "Diffraction-limited imaging. II. Optical aperture synthesis imaging of two binary stars," Astron. J. **97**, 1510–1521 (1989).
12. D. F. Buscher, C. A. Haniff, J. E. Baldwin, and P. J. Warner, "Detection of a bright feature on the surface of Betelgeuse," M. Not. R. Astron. Soc. **245**, 7P–11P (1990).
13. J. G. Robertson, T. R. Bedding, R. G. Marson, P. R. Gillingham, R. H. Frater, J. D. O'Sullivan, and R. P. Norris, "High resolution imaging by optical aperture synthesis—first results from the MAPPIT project," Proc. Astron. Soc. Aust. **9**, 162–163 (1991).
14. W. N. Weir, A. M. Ghez, P. W. Gorham, C. A. Haniff, S. R. Kulkarni, K. Matthews, and G. Neugebauer, "Infrared non-redundant mask imaging at Palomar," in *Amplitude and Intensity Spatial Interferometry*, J. B. Breckinridge, ed., Proc. Soc. Photo-Opt. Instrum. Eng. **1237**, 274–285 (1990).
15. C. A. Haniff, D. F. Buscher, J. C. Christou, and S. T. Ridgway, "Synthetic aperture imaging at infrared wavelengths," M. Not. R. Astron. Soc. **241**, 51P–56P (1989).
16. J. C. Christou, A. Y. S. Cheng, and E. K. Hege, "Seeing calibration of optical astronomical speckle interferometric data," Astron. J. **90**, 2644–2651 (1985).
17. J. C. Christou, "Application of speckle interferometry techniques: working with real data," in *High Resolution Imaging by Interferometry*, F. Merkle, ed., Vol. 29 of European Southern Observatory Conference Proceedings (European Southern Observatory, Garching bei München, Germany, 1988), pp. 97–111.
18. F. Roddier, "The effects of atmospheric turbulence in optical astronomy," Prog. Opt. **19**, 281–376 (1981).
19. J. W. Goodman, *Statistical Optics* (Wiley, New York, 1985).
20. D. F. Buscher, "Optimising a ground-based optical interferometer for sensitivity at low light levels," M. Not. R. Astron. Soc. **235**, 1203–1226 (1988).
21. G. R. Ayers and J. C. Dainty, "Iterative blind deconvolution method and its applications," Opt. Commun. **13**, 547–549 (1988).
22. A. C. S. Readhead, T. Nakajima, T. J. Pearson, G. Neugebauer, J. B. Oke, and W. L. W. Sargent, "Diffraction-limited imaging with ground-based optical telescopes," Astron. J. **95**, 1278–1296 (1988).
23. See, e.g., D. M. Alloin and J.-M. Mariotti, eds., *Diffraction-Limited Imaging with Very Large Telescopes*, NATO ASI Series C (Kluwer, Dordrecht, The Netherlands, 1990), Vol. 274.

24. D. Korff, "Analysis of a method for obtaining near-diffraction-limited information in the presence of atmospheric turbulence," J. Opt. Soc. Am. **63**, 971–980 (1973).

25. C. A. Haniff, D. F. Buscher, J. C. Christou, and S. T. Ridgway, "Diffraction-limited imaging at IR wavelengths using aperture masks and fully-filled apertures," in *Amplitude and Intensity Spatial Interferometry*, J. B. Breckinridge, ed., Proc. Soc. Photo-Opt. Instrum. Eng. **1237**, 259–271 (1990).

26. F. Roddier, "Triple correlation as a phase-closure technique," Opt. Commun. **60**, 145–148 (1986).

27. T. J. Cornwell, "Radio interferometric imaging of weak objects in conditions of poor phase stability—the relationship between speckle masking and phase-closure methods," Astron. Astrophys. **180**, 269–274 (1987).

28. E. Pehlemann and O. von der Lühe, "Technical aspects of the speckle masking phase reconstruction algorithm," Astron. Astrophys. **216**, 337–346 (1989).

29. P. W. Gorham, A. M. Ghez, S. R. Kulkarni, T. Nakajima, G. Neugebauer, J. B. Oke, and T. A. Prince, "Diffraction-limited imaging. III. 30mas closure phase imaging of six binary stars with the Hale 5m telescope," Astron. J. **98**, 1783–1799 (1989).

30. H. Bartelt, A. W. Lohmann, and B. Wirnitzer, "Phase and amplitude recovery from the bispectrum," Appl. Opt. **23**, 3121–3129 (1984).

31. A. Lannes, "On a new class of iterative algorithms for phase-closure imaging and bispectral analysis," in *High Resolution Imaging by Interferometry*, F. Merkle, ed., Vol. 29 of European Southern Observatory Conference Proceedings (European Southern Observatory, Garching bei München, Germany, 1988), pp. 169–180.

32. J. C. Marron, P. P. Sanchez, and R. C. Sullivan, "Unwrapping algorithm for least-squares phase recovery from the modulo 2π bispectrum," J. Opt. Soc. Am. A **7**, 14–20 (1990).

33. C. A. Haniff, "Least-squares Fourier phase estimation from the modulo 2π bispectrum phase," J. Opt. Soc. Am. A **8**, 134–140 (1991).

34. T. J. Pearson and A. C. S. Readhead, "Image formation by self-calibration in radio astronomy," Ann. Rev. Astron. Astrophys. **22**, 97–130 (1984).

35. F. Roddier, "Redundant versus non-redundant beam recombination in an aperture synthesis with coherent optical arrays," J. Opt. Soc. Am. A **4**, 1396–1401 (1987).

36. D. F. Buscher, "Getting the most out of C.O.A.S.T.," Ph.D. dissertation (Cambridge University, Cambridge, 1988).