

Coherent Integrations, Fringe Modeling, and Bootstrapping With the NPOI

Anders M. Jorgensen^a, Dave Mozurkewich^b, Henrique Schmitt^{c,d}, J. Thomas Armstrong^c, G. Charmaine Gilbreath^c, Robert Hindsley^c, Thomas A. Pauls^c, Deane M. Peterson^e

^aLos Alamos National Laboratory, Los Alamos, NM, USA

^bSeabrook Engineering, Seabrook, MD, USA

^cNaval Research Laboratory, Washington, DC, USA

^dInterferometrics, Inc., Herndon, VA, USA

^eStony Brook University, Stony Brook, NY, USA

ABSTRACT

Atmospheric turbulence is a major impediment to ground-based optical interferometry. It causes fringes to move on ms time-scales, forcing very short exposures. Because of the semi-random phase shifts, the traditional approach averages exposure power spectra to build signal-to-noise ratio (SNR). This *incoherent* average has two problems: (1) A bias of correlated noise is introduced which must be subtracted. The smaller the visibility/the fainter the target star, the more difficult bias subtraction becomes. SNR builds only slowly in this case. Unfortunately, these most difficult small visibility baselines contain most of the image information. (2) Baseline phase information is discarded. These are serious challenges to imaging with ground based optical interferometers. But if we were able to determine fringe phase, we could shift and integrate all the short exposures. We would then eliminate the bias problem, improve the SNR, and we would have preserved most of the phase information. This *coherent* averaging becomes possible with multi-spectral measurements. The group delay presents one option for determining phase. A more accurate approach is to use a time-dependent model of the fringe. For the most interesting low-visibility baselines, the atmospheric phase information can be *bootstrapped* from phase determinations on high-visibility baselines using the closure relation. The NPOI, with 32 spectral channels and a bootstrapping configuration, is well-suited for these approaches. We will illustrate how the fringe modeling approach works, compare it to the group-delay approach, and show how these approaches can be used to derive bias-free visibility amplitude *and* phase information. Coherent integration provides the highest signal-to-noise (SNR) improvement precisely in the situations where SNR builds most slowly using incoherent averaging. Coherent integration also produces high-SNR phase measurements which are calibration-free and thus have high real uncertainties as well. In this paper we will show how to coherently integration on NPOI data, and how to use baseline visibilities and calibrate coherently integrated visibility amplitudes.

1. INTRODUCTION

In this paper we are going to show how to coherently integrate in post-processing on data obtained with the Navy Prototype Optical Interferometer¹ (NPOI). Interferometry at visible and near-infrared wavelengths is complicated by the atmosphere. The combined effect of atmospheric turbulence and vibrations in the interferometer causes fringes to shift on short time-scales, sometimes as short as several milliseconds. These shifts are large compared to the fringe spacing, and limits integration times to only a few milliseconds. In that short time it is not possible to obtain a complete measurement on a typically faint star. Therefore, we must record many such short exposures and combine them in post-processing. The conventional approach to this post-processing is the so-called incoherent average, which averages fringe Fourier power. When the fringe contrast is small or the count rate is small, which covers most of the most interesting measurements, a large

bias is introduced which limits the signal to noise ratio achievable. An alternative approach, which has been discussed by several authors in the past²⁻⁶ is coherent integration. In the next section we will describe what coherent integration is, then how to determine the phases needed for coherent integration in section 3. In section 4 we briefly outline why coherent integration is necessary. In section 5 we will discuss some of the issues that need to be overcome when coherently integrating, including the introduction of phase terms, whose solution is presented in section 6, and a reduction of the amplitude of the coherently average, whose solution is presented in section 7. Then we present a section discussing the coherent bias, and finish with a discussion of some implications and observations regarding coherent integration.

2. INTRODUCTION TO COHERENT INTEGRATION

From individual short exposures, we can estimate the complex visibility $X + iY$. On the NPOI, we use a discrete Fourier transform as the estimator. This complex visibility is rotated by atmospheric and other phase variations by some phase $-\theta$ relative to the true source phase. Therefore it is not possible to improve SNR by simply averaging many individual measurements, $X + iY$. The traditional incoherent average bypasses this problem by instead averaging the square modulus,

$$\langle (X + iY)(X + iY)^* \rangle \quad (1)$$

As we shall see, this process makes poor use of the SNR inherently present in the data.

In coherent integration, we instead form the average

$$\langle (X + iY) e^{\theta} \rangle \quad (2)$$

This process preserves far more of the SNR inherently present in the measured data. In order to carry out this process we need however to determine the phase, θ , by which to rotate the individual complex visibilities measurements before adding them.

3. PHASE TRACKING

In order for the coherent integration approach to work, we must have a method of determining the correct phase to apply to rotate each single-frame complex visibility before adding. There are in principle many ways of doing this. The most well-known method for fringe tracking is the group-delay method. In this method the phase as a function of wavelength is assumed to take on the form,

$$\theta(\lambda) = \frac{2\pi d}{\lambda} + \phi \quad (3)$$

where d is a vacuum path delay (the fringe-tracking error), and ϕ represents other phase terms, including a crude approximation of the atmospheric dispersion.

The visibilities at each wavelength are coherently added after rotation by this phase, and the power computed. The maximum of the power is found as a function of d , and the appropriate values to use for coherent integration for that measurement are d , and the negative of the phase of the summed complex quantity. The exact phase reference to use does not matter much, but it is important that the same phase reference

be used for all measurements. The group delay method is only an optimal estimator of fringe phase when (a) the phase is described by equation 3, (b) the count rate and visibility are constant as a function of wavelength, and (c) the wavelength channels are linearly spaced in wave number space. It is also difficult to extend the group delay method to take into account known frame to frame variations due to the slow variation of air-path. For that reason, we have employed a different approach which provides a more generalized approach to fringe tracking. The method was developed for the NPOI, but could be readily adapted to other multi-spectral interferometers.

In the fringe modeling approach we make a model of the fringe,

$$I(\lambda, x) = \sum_{i=1}^N I_i + \sum_{i,j} \sqrt{I_i I_j} \tilde{V}_{ij} \cos\left(\frac{2\pi k_{ij} x}{\lambda} + \theta_{ij}(\lambda)\right) \quad (4)$$

which we fit to the measured fringe pattern. This model correctly takes into account count-rate and visibility variation with wavelength, and can use any appropriate statistics (e.g. Gaussian, Poisson, or any other distribution). By allowing the parameters to vary with time, we can also fit a time-dependent model, which improves SNR. The entire expression is fitted to multiple consecutive frames of data simultaneously, and the best fit tracking phase $\theta_{ij}(\lambda)$ as a function of wavelength is obtained, for example in a form similar to that used in the group delay method, equation 3, or some more realistic variation which also takes into account the dependence of atmospheric dispersion with wavelength, as well as the dispersion in the instrument, and in some cases even phase variations due to the source.

4. IMPORTANCE OF COHERENT INTEGRATION

The simplest reason why coherent integration is important is that it improves SNR. An improvement of the SNR of visibility amplitudes (or V^2), also results in improving the phase, such that single baseline phases can become usable, as we shall see in a later section. The definition of SNR is

$$\text{SNR} = \frac{V^2}{\sigma_{V^2}} \quad (5)$$

Then the incoherent average SNR is⁷

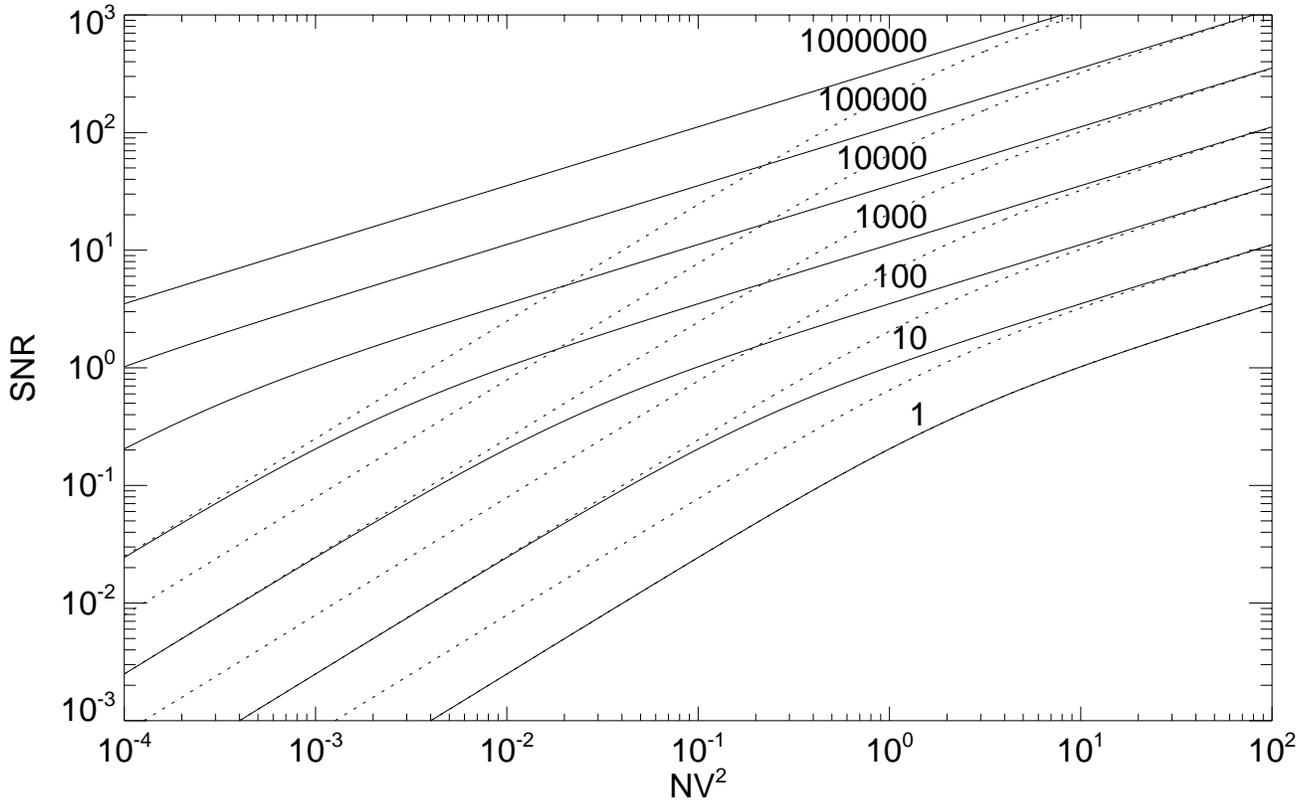


Figure 1. Comparison of incoherent (dotted) and coherent (solid) integration SNR for different 10^0 , 10^1 , 10^2 , 10^3 , 10^4 , and 10^5 combined frames.

$$\text{SNR}_i = \frac{1}{4} \sqrt{M} \frac{NV^2}{\sqrt{1 + \frac{1}{2}NV^2}}. \quad (6)$$

The coherent SNR is

$$\text{SNR}_c = \frac{1}{4} \frac{MNV^2}{\sqrt{1 + \frac{1}{2}MNV^2}} \quad (7)$$

These two functions are plotted in Figure 1. We notice that when the single measurement SNR is high (right side of the plot), we do not gain from coherently integrating. When the single measurement SNR is low on the other hand (left side of the plot), the gains can be dramatic. For example, when $NV^2 = 10^{-3}$, it takes 10^8 measurements to reach a SNR of 1 by the incoherent method, whereas it takes only 10^4 frames to reach the same SNR when coherently integrating. The necessary observing time is reduced by a factor of 10^4 .

There can be little doubt that coherent integration is vital to pushing stellar interferometry to increasingly more detailed measurements requiring pushing the SNR limit of the data.

5. ISSUES WITH COHERENT INTEGRATION

Coherent integration does present some issues which must be dealt with. These include additional phase terms, which, for example, are not present in triple-phases, as well as a reduction of the visibility amplitude due to noise in the determination of the correct tracking phase.

5.1. Extra phase terms

The coherently integrated visibilities, \tilde{V} , contain phase terms in addition to the source phase terms. In the triple-product, these other phase terms are automatically eliminated. However, the coherently integrated baseline visibilities contain phase terms created by both the instrument and by the atmosphere in addition to the source phase terms. There are thus three components to the phase,

$$\phi(\lambda) = \phi_{\text{source}}(\lambda) + \phi_{\text{inst}}(\lambda) + \phi_{\text{atm}}(\lambda) \quad (8)$$

5.2. Phase noise changes the amplitude

Noise in determining the correct phase to rotate the individual frame complex visibilities by will cause a reduction in the resulting coherently integrated visibility. If we consider only phase noise, and not amplitude noise, the visibility reduction factor can be determined as

$$\gamma = \int_{-\infty}^{\infty} \rho(\delta\theta) e^{i\delta\theta} \delta\theta \quad (9)$$

If the phase noise distribution is symmetric, which we assume that it generally is, then γ is a real-valued quantity which reduces the amplitude of the resulting coherent average.

6. SEPARATING EXTRA PHASE TERMS

The phase of the coherently integrated visibility contains three components, as shown in equation 8. The first term is the source phase term, which is the information that we are ultimately after. The second term is an instrumental term, caused by dispersion internal to the instrument. The third term is an atmospheric term caused by dispersion in the atmosphere, and the fringe-tracking loop's attempt at tracking these variations. For most common source geometries the source term can be computed at a parameterized model. The atmospheric phase term, and associated fringe-tracking error can be represented by the following three-parameter model,

$$\phi_{\text{atm}} = \frac{2\pi [(n-1)a + d]}{\lambda} + \phi_0 \quad (10)$$

where a represents the air path mismatch between the two telescope beams forming the baseline, d the fringe-tracking error (a vacuum path), and ϕ_0 a wavelength-independent phase offset.

6.1. Determining instrumental phase

The instrumental phase can be determined by observing a source whose source phase is known to be zero, such as an unresolved calibration star. Figure 2a shows the phase of a calibrator star as the solid curve, and the best-fit atmosphere model, equation 10, as a dashed curve. Note that this is the best fit in a χ^2 sense. If the fit is optimal, then the instrumental phase is the residual, or the difference between the measured phase and the atmosphere model. We repeat this procedure for several more calibrators (20 altogether) observed on one night, average all the residuals to obtain the curve plotted in

Figure 2b. It is always risky to fit a model to data which the model does not perfectly describe, using a χ^2 cost function. In this case the resulting uncertainty of the fit may be a major contributor to the residuals plotted in Figure 2c. In the future we will produce a proper parameterized model for the instrumental phase and fit all terms together to make a better fit and therefore a more accurate determination of the instrumental phase.

6.2. Putting it all together

Once we have a description of the instrumental phase, and a model for the phase of the source, we can combine those with the atmospheric phase (Equation 10), and fit the combined model to a measured phase in order to extract the source parameters. If the instrumental phase is a constant, as we have modeled it thus far, it is simplest to subtract it before fitting the source and atmosphere terms.

6.3. Fitting a binary star

The phase of a binary star relative to its photometric center can be written as

$$\phi_{\text{binary}} = \tan^{-1} \left(\frac{\sin(r\beta) - r \sin(\beta)}{\cos(r\beta) + r \cos(\beta)} \right), \quad (11)$$

where

$$\beta = \frac{2\pi \vec{B} \cdot \vec{s}}{(1+r)\lambda}, \quad (12)$$

and \vec{B} is the baseline vector, \vec{s} is the relative position vector of the two stars composing the binary, and r is the relative brightness of the two stars composing the binary. Figure 3 shows how we separate the source terms from the atmospheric and instrumental terms for a single-baseline measurement on the binary star Mizar.

6.4. Unknown source phase

We can even use this approach in the case of an unknown source phase, or when we do not have a perfect model for the source phase. We will illustrate this with a single-baseline measurement on Vega. In this case we observed Vega on three baselines, two short and one long, and we tracked fringes on the two short baselines to be able to bootstrap and coherently integrate on the long baseline. The long baseline contains a visibility minimum. It is not actually a null, because Vega's asymmetry causes the phase to rotate smoothly from 0 to 180° as a function of wavelength through the

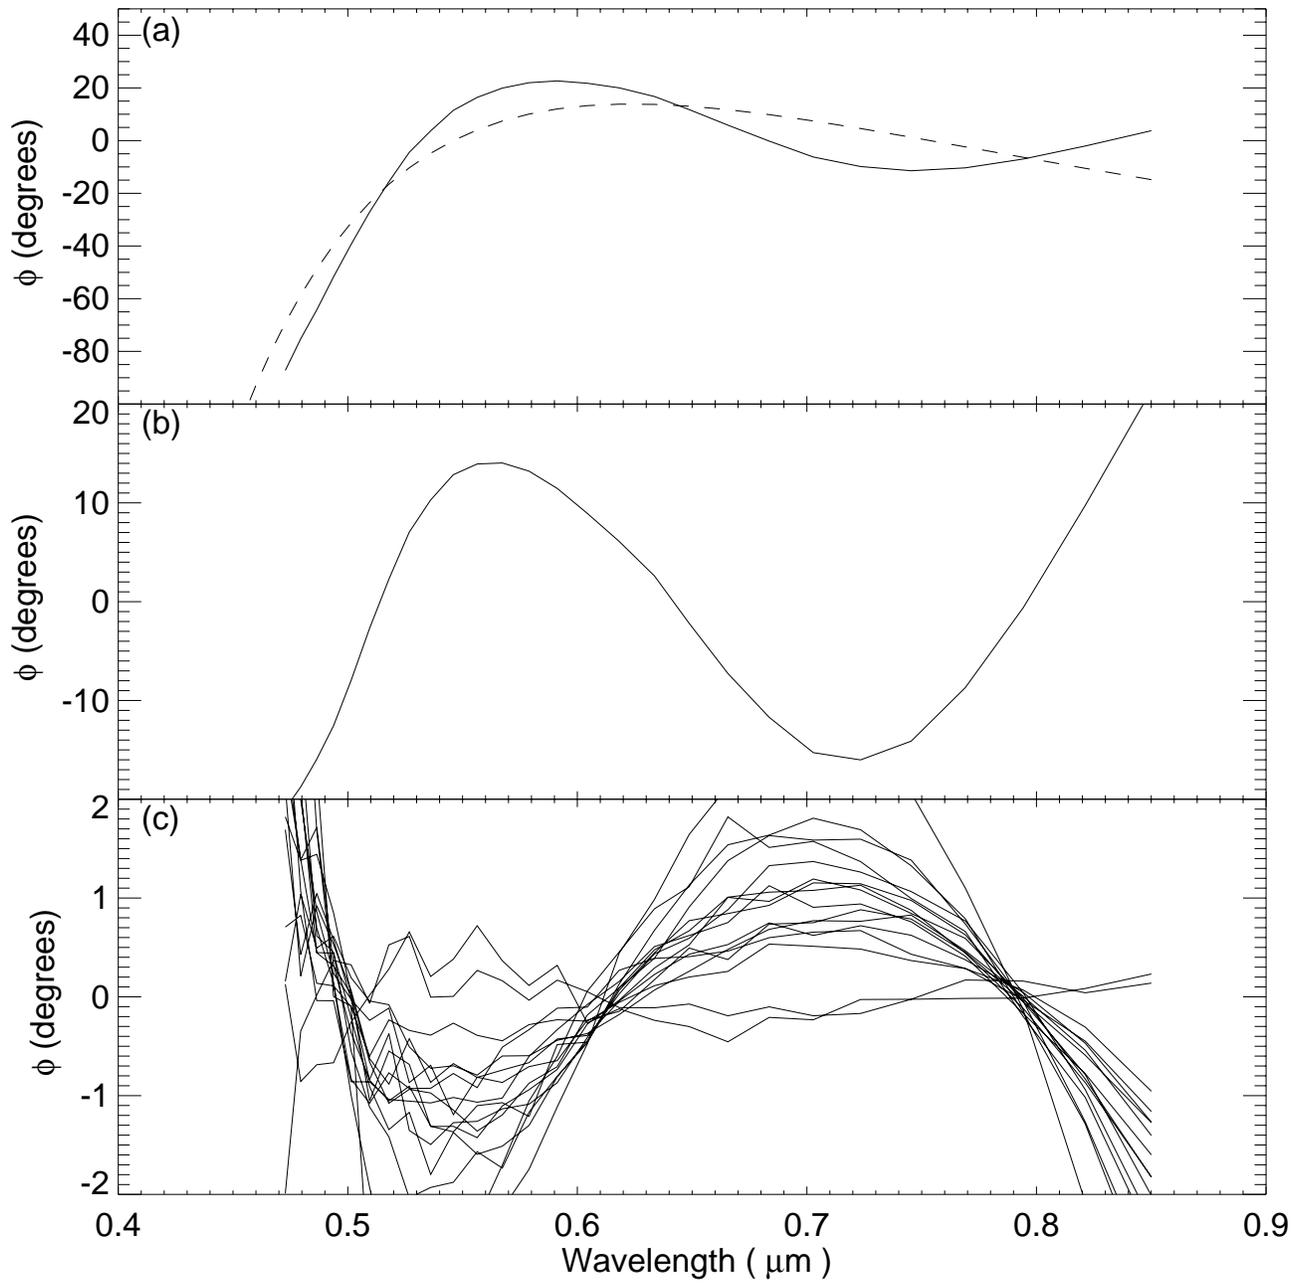


Figure 2. (a) phase as a function of wavelength for a calibration star (solid), and best-fit atmosphere phase model (dashed), (b) average instrumental phase for 20 calibrators, (c) residual for each of 20 calibrators.

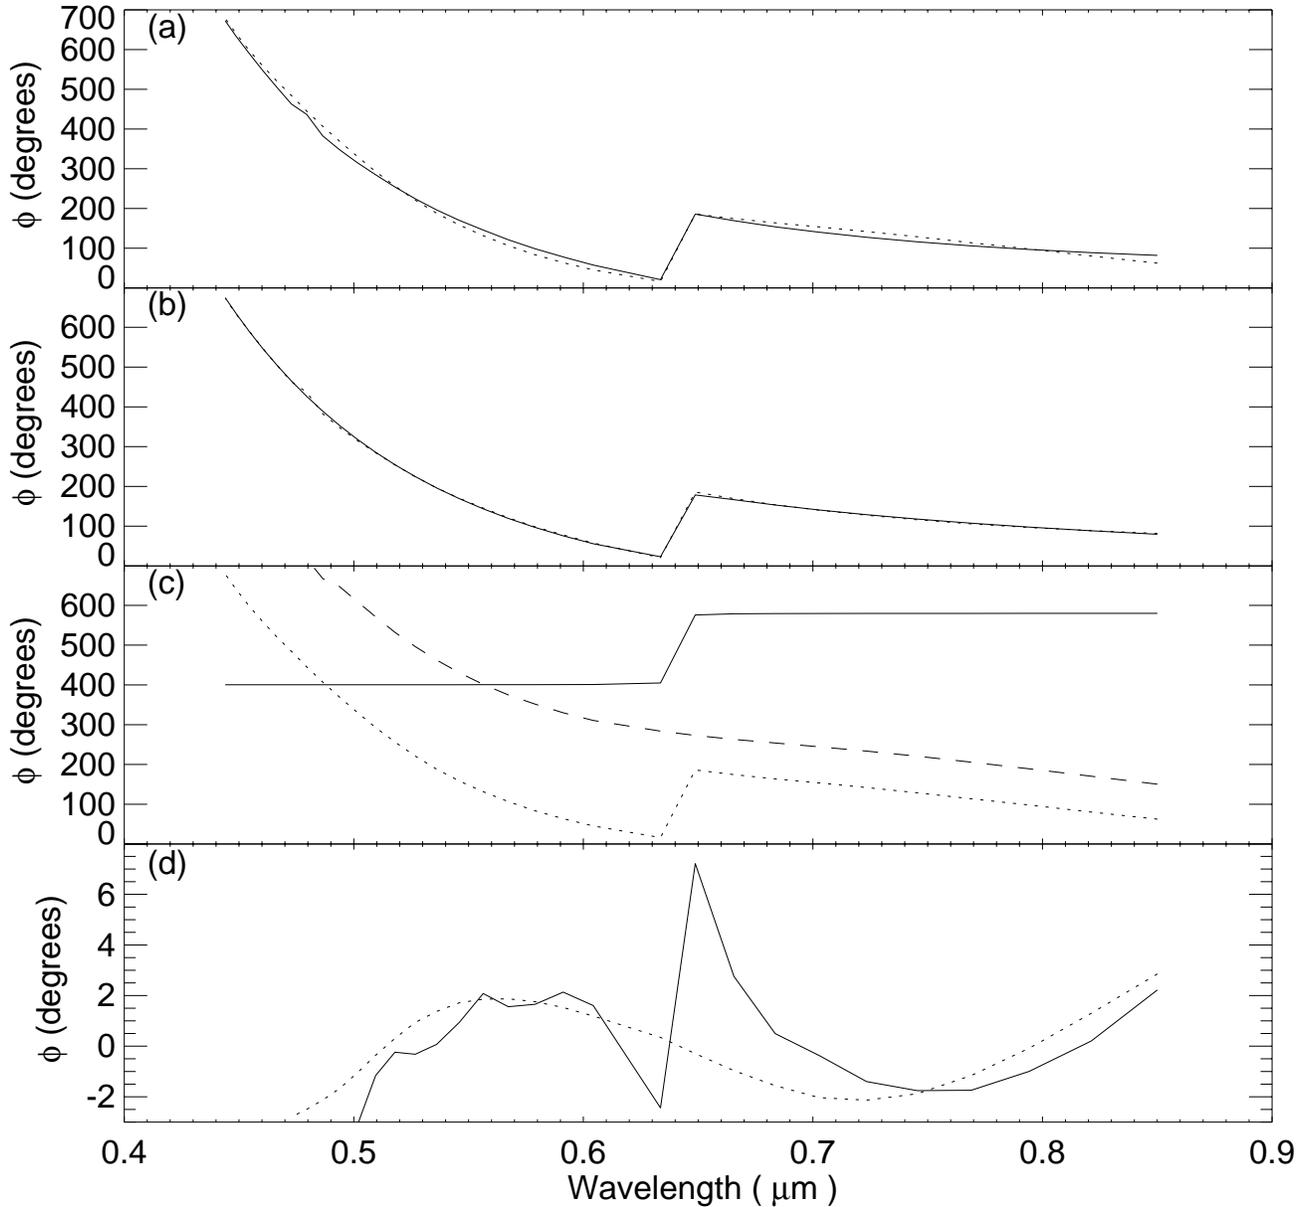


Figure 3. Processing phase information for the binary star Mizar. (a) Measured phase (dotted), and after subtraction of instrumental phase (solid). (b) Source and atmosphere phase (solid) and model fit (dotted). (c) Measured phase (dotted), external phase (atmosphere and instrument) (dashed), source phase (solid). (d) Residual from model fit (solid), scaled instrumental phase (dotted).

minimum. Far from the minimum, the phase is either 0 or 180° . To avoid having to model the unknown phase rotation at intermediate wavelengths, we can create a source model which contains a 180° phase jump at intermediate wavelengths, and fit this model plus the atmosphere to the short and long wavelength data only, after subtracting the instrumental phase term. The result is shown in Figure 4.

7. CALIBRATING AMPLITUDE

When coherently integrating, the phase is only determined to within some uncertainty. This uncertainty, when averaged over many combined measurements results in a reduction of the visibility. If we consider only constant amplitude phasors we can write the coherent integration problem as

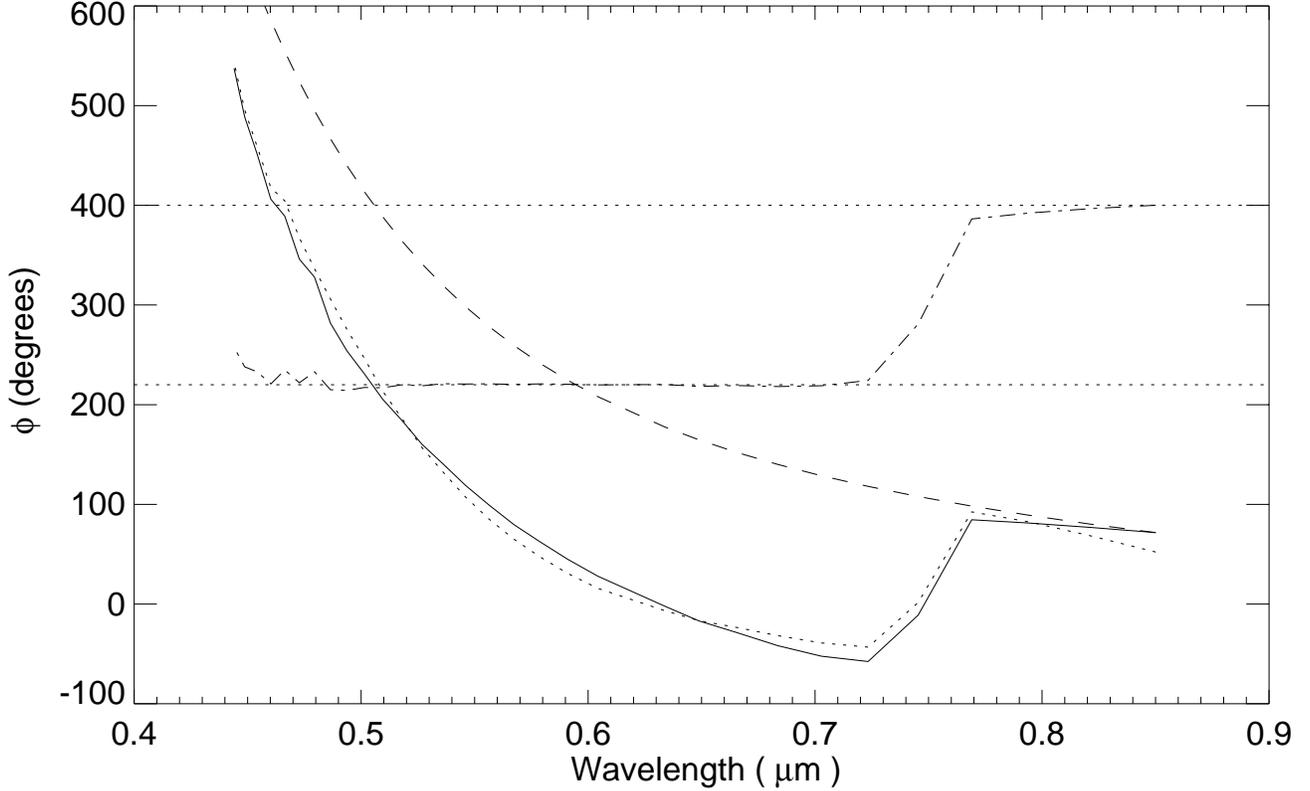


Figure 4. Extracting the Vega source phase by subtracting the instrumental phase and fitting out the atmospheric phase term. The dotted curve is the measured phase. The solid curve is the combined atmospheric and source phase (after subtracting the instrumental phase from the measured phase). The dashed curve represents the atmospheric phase term. The dash-dotted curve represents the source phase.

$$\gamma = \int_{-\infty}^{\infty} \rho(\delta\theta) e^{i\delta\theta} d\delta\theta \approx \sum_{i=1}^N e^{i\delta\theta_i} \quad (13)$$

If ρ is a zero-mean Gaussian with standard deviation σ , then the integral evaluates to

$$\gamma = e^{-\sigma^2}, \quad (14)$$

where σ is the phase noise. We can determine the phase noise from the ratio of the incoherent and coherent squared visibilities,

$$\sigma = \sqrt{\log \left(\sqrt{\frac{\langle V^2 \rangle}{\langle V \rangle^2}} \right)}, \quad (15)$$

where $\langle V^2 \rangle$ is the incoherent average, and $\langle V \rangle^2$ is the coherent average. When bootstrapping baselines, the phase noise adds in the usual way,

$$\sigma_3^2 = \sigma_1^2 + \sigma_2^2 \quad (16)$$

and the coherently integrated squared visibility can be calibrated as

$$V_{\text{calibrated}}^2 = V_{\text{bootstrapped}}^2 e^{4\sigma_3^2} \quad (17)$$

7.1. Example: Alcaid

We will first illustrate this with a calibration star, Alcaid. The reason for showing this example is that we can obtain high-SNR incoherently averaged V^2 on all baselines so that we know the answer that we are trying to obtain. Figure 5 shows the three baselines observed on Alcaid in one observation. In all three panels the incoherent average is shown as a dotted curve, obtained through a simultaneous fit of the fringe power and bias separately for each wavelength (see Mozurkewich⁸ for further details). The uncertainties in those fit are marked as vertical bars. In the first two panels, the solid curve represents the coherently averaged squared

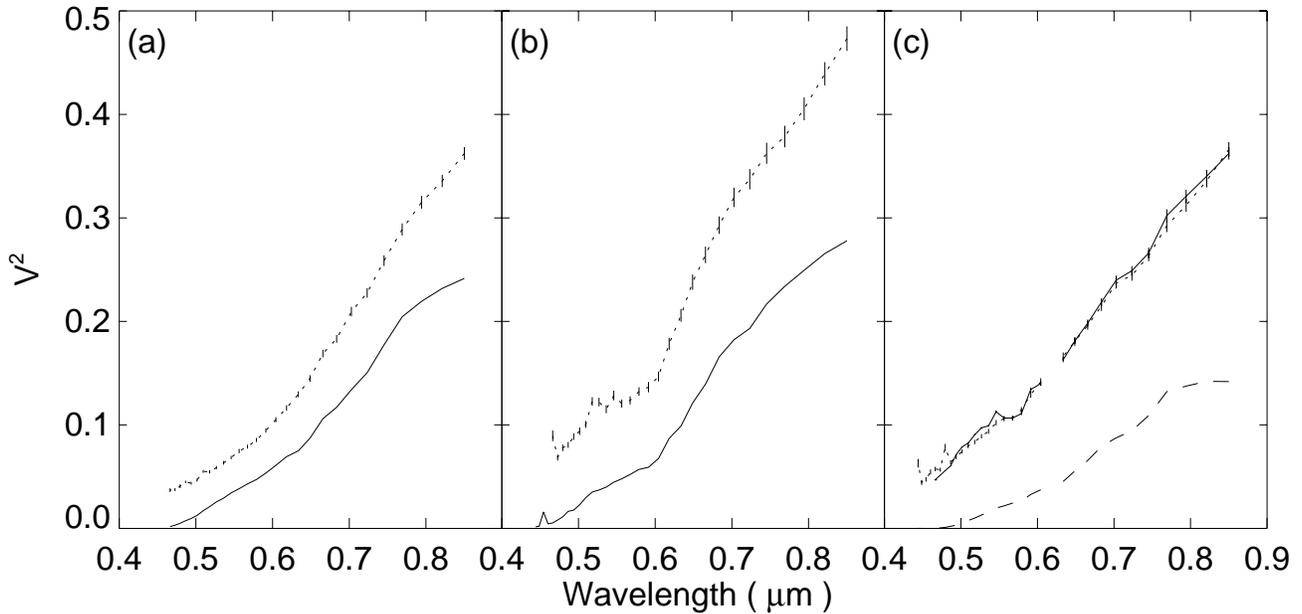


Figure 5. Illustration of the calibration of bootstrapped coherently integrated visibilities, using three unresolved baselines forming a closure triangle on the star Alcaid. Panels (a) and (b) show the tracking baselines, while panel (c) shows the bootstrapped baseline.

visibility (that is, the square-modulus of the coherently averaged visibility). Because of phase noise, the coherently averaged squared visibility is smaller than the incoherently averaged squared visibility, but the SNR on the coherently integrated squared visibility based on photon counting statistics is much better. In the third panel, we bootstrapped the phase-tracking information from the first two panels. The dashed curve is the coherently averaged squared visibility. The solid curve is the corrected coherently averaged squared visibility, obtained using equation 17. The nature of the factor 4 in equation 17 is still not clear. It was obtained experimentally using both real and simulated data. It may be that amplitude noise and not just phase noise, as assumed in equation 13 contributes the additional factor. In that case, we may have to modify the definition of σ in the future.

7.2. Example: Vega

It is far more interesting to apply coherent integration to a situation in which a low-visibility baseline can be bootstrapped by two high visibility baselines, and sufficient SNR on the low-visibility baseline cannot be obtained using the incoherent integration method. We illustrate this case with an NPOI observation of Vega. Figure 6 shows the results of bootstrapping and coherently integrating on a baseline which resolves Vega. Figure 6 corresponds to panel c in figure 5. As before, the dotted curve is the incoherent integration, the dashed

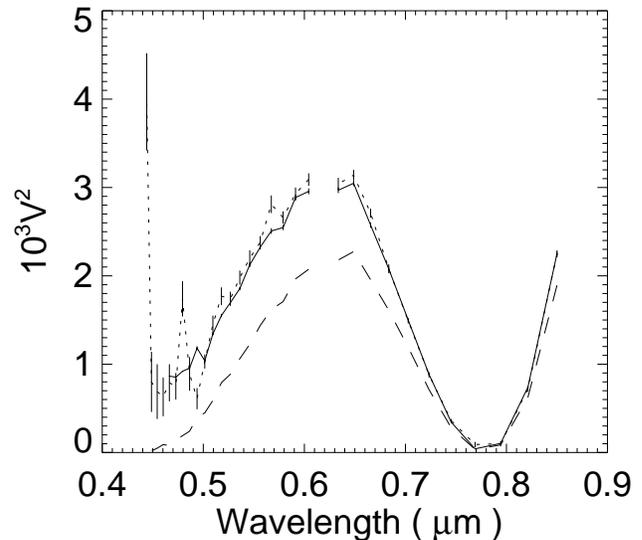


Figure 6. Bootstrapping and coherent integration on a Vega baseline: the dotted curve shows the incoherent integration, the dashed curve the bootstrapped coherent integration, and the solid curve the corrected bootstrapped coherent integration.

curve is the bootstrapped coherent integration, and the solid curve is the corrected bootstrapped coherent integration. Notice that the visibility reduction factor (the solid curve divided by the dotted curve) is not as small as in Figure 5, because Vega is much brighter than Alcaid, thus having better single-frame SNR, and thus better phase tracking. It is the phase noise on the track-

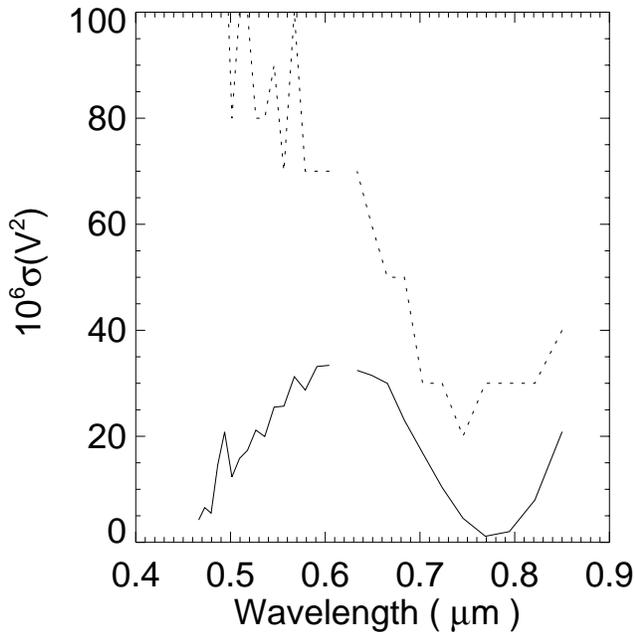


Figure 7. Bootstrapping and coherent integration on a Vega baseline: the dotted curve shows the standard uncertainty on the incoherently integrated V^2 , and the solid curve shows the standard uncertainty on the coherently integrated V^2 .

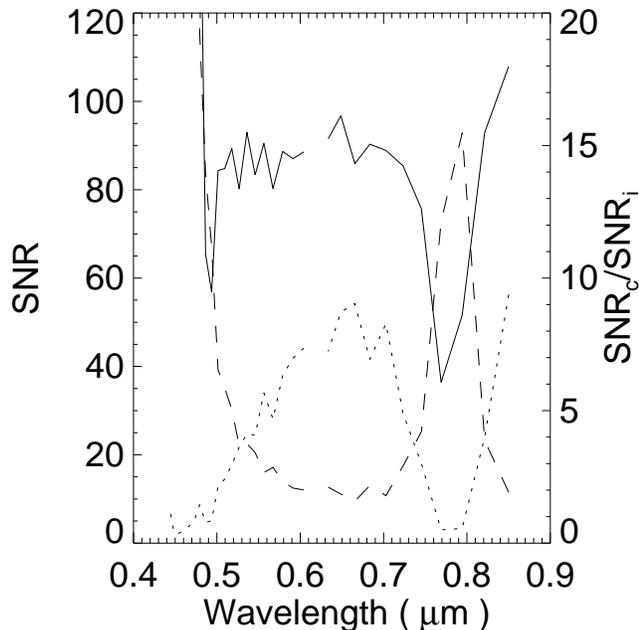


Figure 8. Bootstrapping and coherent integration on a Vega baseline: the dotted curve shows the SNR on the incoherently integrated V^2 , and the solid curve shows the SNR on the coherently integrated V^2 .

ing baselines which determines the visibility reduction on the bootstrapped baseline, which is one of the great advantages of bootstrapping. This is a very important point which we will go into further, in the following section.

7.3. SNR of coherently integrated visibilities

Figure 7 shows the standard deviation of the squared visibilities derived both through incoherent integration (dotted) curve, and through bootstrapping and coherent integration (solid curve). We immediately notice that the two curves look very different, with the standard deviation of the incoherent integration not going through the range of variation of the standard deviation of the coherent integration. The minimum value of the incoherent standard deviation is determined by the photon counting uncertainty on the bias, whereas the value of the standard deviation for the coherent average is primarily determined by the standard deviation of the incoherent V^2 on the tracking baselines, at least until the V^2 on the bootstrapped baseline drops to the photon noise limit. And in this particular case, the photon noise limit is only reached at the smallest visibilities in the plot.

Figure 8 shows the same picture, but this time in more familiar SNR terms. As before, the dotted curve shows the incoherent average SNR, whereas the solid curve shows the coherent average SNR. We have added a third curve, which is the SNR in the coherent average divided by the SNR in the incoherent average, plotted as a dashed curve. The picture is reversed from Figure 7, with the incoherent average SNR being smaller and having a similar variation to V^2 , and the coherent SNR being large and largely independent of V^2 , except for the smallest values of V^2 . The reason is that the SNR in the incoherent case is dependent on the ratio of V^2 to the bias, whereas in the coherent average case the SNR is simply the propagation of the uncertainties in determining the incoherent averages on the tracking baselines, and thus the phase noise, plus the photon noise on the coherent average. It is only when V^2 drops very far that the photon noise begins to contribute. The SNR of the coherent average is thus independent of V^2 as long as V^2 is large enough that there are enough photons in the coherent average. That is the case here except at the bottom of the V^2 minimum between 0.75 μm , and 0.8 μm .

When the amplitude SNR is high, then the phase SNR is also high, and conversely, when the amplitude SNR is low, the amount of information available about

the phase is also low. The improvement in very low SNR situations is therefore not just incremental, it is dramatic, precisely as illustrated in Figure 1.

Let us look again at Figure 8. This plot represents 88 s of data. If we wanted to increase the SNR on the incoherent average by a factor of 16 to reach the SNR of the coherent average, we would, according to equation 6, need to increase observing time by a factor of 256 to 6.25 hours. Using coherent integration therefore also improves observing efficiency, in this case by a factor of 256. Next, let us suppose that we want to improve the SNR to 100 in both the coherent and in the incoherent case. For the coherent case we would need to increase the observing time by a factor of 7.72 to 679 s. In the incoherent case, we would need to increase the observing time by a factor of 1.1×10^3 to 27 hours. Of course, even if we had many hours or days to observe a given star, we would not be able to incoherently integrate those data in the usual way because Earth rotation would change the baselines during that time interval.

But it may not be necessary to integrate for long periods of time. Most observations today do have very large inherent SNR, probably enough to perform most of the analysis that we wish to perform. But it takes coherent integration to fully exploit this inherent SNR.

8. COHERENT BIAS

When tracking fringes, by any method, the fringe tracking is susceptible to fitting noise. Coherently integrating on the same photons which are used for fringe tracking will tend to increase the visibility amplitude.

We have not yet derived an expression for the coherent bias based on the count rate and visibility in the different data channels, but Meisner² suggested an alternative solution. Fringe track on only some of the channels and use that information to coherently integrate the other channels. Meisner² specifically suggested tracking fringes on three quarters of the channels and using that tracking information to coherently integrate on the remaining quarter of the channels, and then repeating this procedure 4 times. This procedure does produce an unbiased visibility amplitude, but it introduces two new problems. First, we are no longer using all of the available information to track fringes. This will increase the phase noise, and produce coherent amplitudes with smaller SNR. A second problem is that the average remaining atmosphere and fringe tracking error will be different for each quarter of the data, producing phases with this additional structure. This in turn means that a separate atmosphere term (consisting of three parameters) must be removed from

each quarter of the data. These two problems speak for solving the problem of determining the coherent bias and subtracting it. We have found experimentally that using this approach to measure the visibility reducing phase noise works well because the periodic variation in the atmosphere and fringe tracking error terms are small relative to the incoherent average SNR.

For the phases on the other hand, care must be taken. Slightly different residual air and fringe tracking error paths in the fringe tracking of the different sets of channels will produce slightly different phases in the different sets of channels, which must be treated appropriately. If we have four sets of channels, we thus have four times as many atmosphere parameters.

If we take the approach of maximizing the amount of information that goes into fringe tracking, we can repeat the fringe tracking once for every channel, leaving out only that channel, and tracking on all other channels. In the case of NPOI, which has 32 channels, this would increase the computational burden by a factor of 31. This may be reasonable for bootstrapping and amplitude calibration in cases where the SNR is very low. When we use phases, this approach may not work, because we now have a situation in which we have a different atmosphere term for each channel, and insufficient information to fit for all of those terms. It is possible that we can remedy this by using information about the actual fringe tracking terms that were derived for each channel, and include the average phase relative to a reference channel as part of the description of the atmosphere term.

9. DISCUSSION

Baselines phase has better SNR than triple phase. Baseline phases also provide more information than triple phases in some cases of complex source geometry. Phases are calibration-free. They are not biased by atmospheric conditions in the same way that amplitudes are. This means that we can take full advantage of the formal SNR in the measurement, which in many cases is a small fraction of a degree (for example, in Figure 3) the formal error is 0.03° at the red end of the spectrum, and 0.2° at the blue end of the spectrum). Visibility amplitudes are biased to smaller amplitudes when coherently integrating, due to phase tracking noise. When bootstrapping we can measure the phase tracking noise on the tracking baselines by comparing them to the incoherently averaged V^2 on the same baselines, and bootstrap the phase noise on the bootstrapped baseline. The visibility on the bootstrapped baseline can then be corrected for this reduction.

Coherent integration is part of the necessary future advancement in stellar interferometry. It provides better SNR, in some cases dramatically so, exactly in the situation where the incoherent integration method fails to produce a good SNR. These situations are usually cases of low visibility because a baseline fully resolves the target object, or cases in which the target object is very faint. In the first case the measurement is interesting because it is one which provides information about the structure of the target object. In the second case the measurement is interesting because it provides better SNR on faint sources and thus opens the possibility of observing a greater variety of faint sources.

10. CONCLUSION

The takeaway messages of this paper are

1. We can coherently integrate to produce high SNR complex visibilities.
2. We know how to deal with the additional terms in the baseline phase.
3. We know how to calibrate the visibility amplitude when bootstrapping.
4. Phases are calibration free, which means that we can make the true uncertainty equal to the formal (photon noise) uncertainty.
5. Bootstrapping and coherent integration provides high SNR even when the visibility is very small, which are the most interesting cases, and the cases where the incoherent method fails.
6. Low SNR measurements have much higher SNR than we are able to extract using incoherent integration. However, it takes coherent integration to fully take advantage of the SNR in those cases.
7. Future interferometers should be built with spectral capability and the ability to observe at least three baselines simultaneously. This allows both coherent integration and bootstrapping.

ACKNOWLEDGMENTS

The NPOI is funded by the Office of Naval Research and the Oceanographer of the Navy. This work was also supported by the U.S. Department of Energy and the University of California.

REFERENCES

1. J. T. Armstrong, D. Mozurkewich, L. J. Rickard, D. J. Hutter, J. A. Benson, P. F. Bowers, N. M. Elias II, C. A. Hummel, K. J. Johnston, D. F. Buscher, J. H. Clark III, L. Ha, L.-C. Ling, N. M. White, and R. S. Simon, "The Navy Prototype Optical Interferometer," *Astrophys. J.* **496**, pp. 550–571, 1998.
2. J. Meisner, "Coherent integration of fringe visibility employing probabilistic determinations of atmospheric delay," **3350**, pp. 294–308, 1998.
3. C. A. Hummel, D. Mozurkewich, J. A. Benson, and M. Wittkowski, "Coherent integration using phase bootstrapping," in *Proc. SPIE*, **4838**, pp. 1107–1114, 2003.
4. C. A. Hummel, D. Mozurkewich, J. A. Benson, and D. M. Peterson, "Coherent integration of npoi phase closure data on altair," in *Proc. SPIE*, **5491**, pp. 707–714, 2004.
5. A. M. Jorgensen, D. Mozurkewich, T. Armstrong, R. Hindsley, T. Pauls, C. Gilbreath, and S. Restaino, "Fringe fitting for coherent integrations with the npoi," **5491**, pp. 1474–1485, 2004.
6. J. Meisner, R. Tubbs, and W. Jaffe, "Coherent integration of complex fringe visibility employing dispersion tracking," **5491**, pp. 725–740, 2004.
7. M. Shao, M. M. Colavita, B. E. Hines, D. H. Staelin, D. J. Hutter, K. J. Johnston, D. Mozurkewich, R. S. Simon, J. L. Hershey, J. A. Hughes, and G. H. Kaplan, "The mark III stellar interferometer," *Astron. Astrophys* **193**, pp. 357–371, 1988.
8. D. Mozurkewich, J. T. Armstrong, G. C. Gilbreath, and T. A. Pauls, "Estimation of fringe parameters," **5491**, pp. 753–762, 2004.