Coherent integration of fringe visibility: a generalized approach

Jeffrey A. Meisner

Leiden Observatory, Postbus 9513, 2300 RA Leiden, The Netherlands

ABSTRACT

The measurement of fringe visibility (correlation) in a long baseline optical/infrared astronomical interferometer may be performed through a coherent integration of noisy data if a measurement of the phase producing the interference is available. That phase reference can be applied in hardware by rapidly servoing a delay line, or may be applied offline to a data stream produced while the phase was being monitored. In a spectrally dispersed detection system, the same data may be used for atmospheric delay tracking as well as visibility estimation. In that configuration one measures not only the magnitude, but the phase of the visibility. The squared magnitude of the measured interference need never be formed, so that all averaging is performed on the unbiased measurements of complex correlation, optimally reducing the effect of measurement noise. However accurate results depend on careful modeling of the system and correction for a number of effects which are detailed in this paper. The results, supported by simulations, are applicable to visibility determinations in a wide range of astronomical interferometers utilizing analog detection or photon counting. In particular these methods are being developed and proposed for use with the MIDI 10 micron interferometric instrument being built for the VLTI.

Keywords: Astronomical interferometry, fringe visibility, fringe tracking, coherent integration, estimation, phase referencing

1. INTRODUCTION

As opposed to incoherent estimation techniques traditionally used in astronomical interferometry, there has been increased interest in estimators based on *coherent*[1] or *phase-referenced*[2–4] estimation of optical correlation (or "fringe visibility"). These related and overlapping concepts, involve the *explicit* use of a determination of the differential optical path delay (OPD) $\tau(t)$ in the calculation of the correlation from interferometric data, unlike incoherent estimators which treat the detection of interference essentially as a noise source with no known phase. Broadly speaking, coherent integration can be performed in one of two ways. Using fringe-tracking *hardware* the differential OPD affecting the phase of interference is actively controlled to within a fraction of a wavelength so that the phase of interference is locked for a period of time during which data is thereby acquired coherently. Alternatively, it is possible for data to be acquired during a period in which the differential OPD may be changing substantially, but during which a *second* interferometric detector records changes in the phase of interference. That second data stream is then used to interpret the primary data in order to estimate the amplitude and phase of the optical correlation that generated it. While both approaches have their advantages, their essential capabilities and performances do not differ. In particular, these coherent methods will provide a higher net signal to noise ratio in the visibility measurement compared to incoherent integration for observations of weaker sources, and they will be capable of measuring the phase of interference with respect to the phase reference.

Of course all interferometry is coherent, by definition, for it is only the correlation in the electric fields of the interfering signals that causes the intensity of the combined signal to vary from the sum of the intensities of the individual contributions. However when bandlimited light is interfered, what is detected is only the real part of the net complex correlation. Since the phase of interference will be randomized by atmospheric turbulence, the measured correlation will be the magnitude of the underlying correlation multiplied by the cosine of a random phase. The cosine of a random phase will, of course, average to zero, so if the interference is simply accumulated (averaged) over a long period, the signal will be cancelled.

Thus if one desires to perform long optical integrations in order to overcome detector noise, it is required that the phase affecting the interference be stabilized in hardware. On the other hand, when using detectors which are limited by photon shot noise even during short integrations, then this requirement vanishes, for it is now possible to compensate for the interference phase *following* the detection of the data. In fact, all else being equal, somewhat

Send correspondence to meisner@strw.leidenuniv.nl

superior results can generally be obtained in this manner, by using an improved estimator for the interference phase, one that employs *future* as well as past data.[5]

Use of offline correction for the phase of interference does not necessarily relieve us of all delay tracking requirements, since the detected amplitude of interference can be substantially degraded if $\Delta\nu\,\tau$ is not small compared to unity, where τ is the uncorrected differential OPD and $\Delta\nu$ is the spectral width of the detection channel. Thus we are still required to compensate *in hardware* for large excursions of the atmospheric OPD τ_e, and/or require a narrow spectral width $\Delta\nu$ in the detection channel.

1.1. Sources for the phase reference

Let us briefly look at various sources for the phase reference which may be employed to perform coherent integration of interferometric data. A simple source of a phase reference for data produced interferometrically at optical frequency ν_1 is from tracking the phase of interferometric data received at optical frequency ν_{ref}.[4] In the extreme case of spectrally dispersed detection (corresponding to the simulations presented in this paper) where there are a multitude of independent spectral channels through which interference is detected, this implies using the entire data set to estimate the atmospheric OPD,[5] and that determination is used to create a reference phase to coherently integrate the same data received by each of those detector channels. Especially at lower signal to noise ratios, the dual use of the same data contributes a significant positive bias in the estimation of correlation[6] which can be removed by tailoring the specific phase reference for channel i to exclude from consideration data from that particular detector channel.

At visible wavelengths use of a phase reference at one wavelength will not be fully correlated with the phase affecting interference at a different wavelength due to random and non-random components of material dispersion (designated as ψ_2 and ψ_1 respectively in Figure 2). The instrumental sources of dispersion (for instance, due to the air in the delay line or other unequal optical paths) and the dispersion due to the fluctuations in air density[7] responsible for the first order component of atmospheric path delay variations observed by the fringe tracker, can in principle be calibrated and coherence restored.[3] The random components of dispersion, if not much smaller than a wavelength, would need to be tracked separately.

Use of a phase reference at widely separated wavelengths introduces various opportunities and complications. On the one hand, use of lower frequency (longer wavelength) data for fringe tracking will allow measurement of a small degree of interference at a higher frequency where the same object may be highly resolved.[4] One potential drawback of this method if used over a very wide wavelength range, is that the fringe tracking error will be magnified by the wavelength ratio. Thus if one were to use a fringe tracker at $2.2\mu m$ which had a .5 radian RMS error to reference interferometry at 550nm, the resulting 2 radian error would reduce the measurement (and SNR) to 13% of its full value, and make calibration very imprecise. On the other hand using a fringe tracker at a short wavelength as a phase reference for interferometry at a longer wavelength will not have this problem at all (the phase tracking error will actually be reduced) but may be largely pointless inasmuch as an object that is not overresolved at the tracking wavelength will tend to be underresolved at the longer measurement wavelength. But in some cases additional object structure only emitting at the longer wavelength (e.g. dust shells observable in the infrared around a visible star) would nevertheless call for such an observational mode.

A source of phase referencing which has received a lot of attention in recent times[8,9] involves the use of a nearby *off-axis* astronomical source using a “dual feed” instrument. If the two objects are well within the isoplanatic patch, then the atmospheric phase variations will be sufficiently correlated in order for a phase reference from the brighter object to be used to coherently integrate the light even from a dim object. The additional atmospheric phase noise affecting the light from the untracked object will degrade the magnitude of interference and make calibration of visibilities difficult. Nevertheless this will enable interferometric measurements on very dim objects located at fortuitous positions in the sky. Measurements of the science object’s phase with respect to the reference will not be biased (since the relative atmospheric delay is zero mean), and an instrument whose internal path lengths are carefully monitored at the nanometer level is expected to allow for narrow angle astrometry on micro-arcsecond levels,[10] given one object which is sufficiently bright on which to fringe track.

Another source of phase referencing that has been demonstrated involves the use of the combined phase inferred from fringe tracking on two short baselines to coherently integrate the interference observed using the longer total baseline.[11] This has been successful in measurement of the small visibilities from the overresolved interferometry over the long baseline, yielding measurements which would not have otherwise been possible using direct fringe tracking.

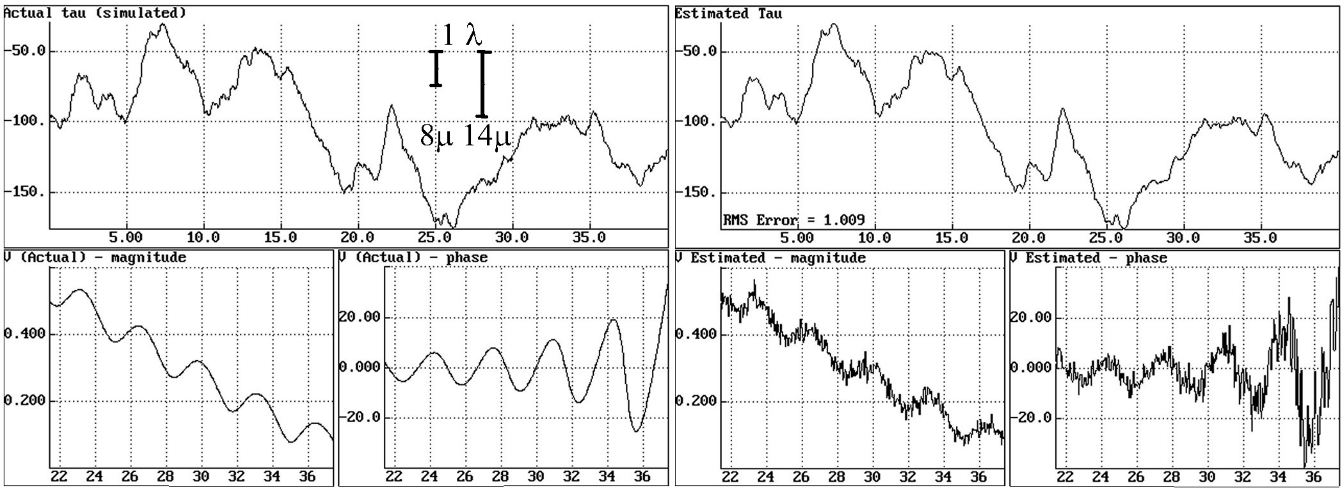

Figure 1. A random atmospheric delay function (upper left) is simulated and plotted in femtoseconds over a period of 40 seconds. Simulated data is thereby produced corresponding to the hardware model of the MIDI instrument with a photometric SNR of 2.0 in each of 640 detector channels. The visibility magnitude and phase in degrees (lower left) is plotted as a function of optical frequency in Terahertz. This corresponds to a hypothetical binary star (see text). The simulated data is used to produce estimates of the OPD (upper right) which is itself used as a phase reference to generate the visibility curve (lower right).

Extension of this method to combining three or more fringe tracked baselines could permit coherent measurements at baselines at which the visibility has fallen to a very small value.

Finally we might mention the use of a phase reference used in coherent integration at the Palomar Testbed Interferometer (PTI). In this case visibility measurements use light from one output of the photomixer which is spatially filtered with a pinhole and spectrally dispersed, both of which substantially reduce the signal to noise ratio. However data from those "science" measurements can be integrated coherently using a phase reference from the light from the opposite output of the photomixer which is neither spatially filtered nor dispersed spectrally, providing a strong signal suitable for fringe tracking on dimmer objects.[1]

1.2. The MIDI instrument for the VLTI

The estimators for V herein and the associated fringe tracking algorithms[5] are proposed for use in the MIDI instrument,[12] the mid-infrared interferometric detector planned for use as part of the VLTI starting in late 2001. MIDI will operate in the 8 - 14 micron band. In one operational mode, the light received from two 8 meter telescopes will be spatially filtered by pinholes, interfered, and spectrally dispersed to produce two spectra each of which is detected across 320 detector pixels. The detector will be read out every 50 milliseconds.

We have simulated such data having a photometric signal to noise ratio of 2.0 in each of the 640 detector channels; this might be obtained using MIDI for an object as dim as magnitude N=3.5. In this simulation we assume a source model consisting of a primary star whose uniform disc has a diameter of 30 femtoseconds (about 14 milli-arc seconds over a 130 meter baseline), and a companion with a relative luminosity of 5% at a separation of 300 femtoseconds. The resulting visibility magnitude and phase (in degrees) versus optical frequency from 22 to 37 Terahertz ($\lambda = 8 - 14 \mu m$) are plotted in the left pane of Figure 1 along with a random differential atmospheric OPD (in femtoseconds) over a period of 40 seconds. Simulated data were produced for each of the 640 detector channels for every 50 millisecond period over a 40 second observation. No intentional delay was introduced beyond compensation for the geometric delay corresponding to the earth's rotation.

Using the atmospheric delay estimation techniques[5] proposed, a good estimate of the OPD is obtained in the right pane of Figure 1, with an RMS estimation error of about 1 femtosecond. That estimate is used as a phase reference to coherently integrate the simulated data obtained over the 40 second observation, and yield an estimation of the complex visibility versus optical frequency whose magnitude and phase are graphed in the right pane of Figure 1. The results appear consistent with the expected estimation error of about .013 RMS in both the real and imaginary parts of V.

2. COHERENT INTEGRATION: THE BASIC APPROACH

Given a system in which a hardware "fringe tracker" has successfully removed all variation in the OPD, then coherent integration is fairly trivial, for any measurement of interfered light will be coherent by definition. But in general a perfect fringe tracker is not possible, so there will be an uncorrected component of atmospheric OPD. We may also choose not even to attempt to track the OPD in hardware, or to add an intentional delay modulation to the optical path. So suppose there is a large or small component of uncorrected OPD $\tau(t)$ for which we have an estimate (we will initially ignore the presence of estimation error). Let us look at the method by which this knowledge can be applied to a data set to produce an accurate estimate of complex fringe visibility V. We will start with a simple model, and briefly review the basis in probability theory for the maximum likelihood estimator in Section 2.1 (which the casual reader may choose to skip). From that rigorous treatment of the estimation problem, we can derive the basic method of coherent integration using a phase reference offline. Following this brief presentation, we will consider a more detailed hardware model and explore nuances of the application of such an estimator, and modifications or corrections required to optimize its performance.

So let us assume we have a set of K interferometric data x_k observed at times t_k. The net differential atmospheric and instrumental OPD at time t_k we shall call τ_k. For simplicity we shall assume our data x_k are from narrowband detection at optical frequency ν from interference at a mixing phase of 0, i.e. the electric fields were added in phase. Further let us assume that we have a good estimate of the photometric signal level s, and that the generated interference has a relative visibility V, a real number which we seek to estimate. We could then model the measured data x_k formed at time t_k as:

$$x_k = s\,(1 + V cos(2\pi\nu\tau_k)) + (\text{noise}) \ . \tag{1}$$

Let us take the noise term to be zero-mean gaussian with a RMS level equal to n.

2.1. The likelihood function

The *a priori* probability density for x_k would clearly be given by:

$$P(x_k) = \frac{1}{\sqrt{2\pi}\,n} \ e^{-\frac{1}{2\,n^2}\ (x_k - s(1+V\,cos(2\pi\nu\tau_k)))^2} \ . \tag{2}$$

Since the noise terms are independent given the above model, the joint probability for the entire data set would be the product of (2) over all k. Once the data has been produced, a reference to the *a priori* probability that would have produced it *now viewed as a function of the model parameters* is known as the *likelihood*. The logarithm of the likelihood, λ, would thus be given by:

$$\lambda = log(P(x_1, x_2, \dots x_K)) = -K\,log(\sqrt{2\pi}\,n) \ - \ \frac{1}{2\,n^2}\sum_k (x_k - s(1+V\,cos(2\pi\nu\tau_k)))^2 \ . \tag{3}$$

From the likelihood function produced by the data set x_k, there are two essentially orthogonal estimations that can be performed. First, a good estimate of the random atmospheric OPD $\tau(t)$ can be obtained using Bayesian estimation. In that case the *a posteriori* probability density assigned to a function $\tau(t)$ is found by multiplying the likelihood e^λ by the *a priori* probability density for that particular delay function. Maximization of that quantity will yield the *Maximum A Posteriori Probability* (MAP) estimator for the delay function. Although an assumed value for the fringe visibility V must be supplied, the resulting estimate of $\tau(t)$ has no first order dependence on the magnitude of that assumed parameter. This method of atmospheric OPD estimation has been previously detailed[5] and will therefore not be further described in this paper. However it has been used in many of the simulations to be presented, and is considered an essential part of the method, except when an alternate source of fringe tracking is available.

The second quantity that can be estimated is, of course, the fringe visibility V. In this case we require knowledge of the atmospheric OPD function $\tau(t)$, but again, the estimation of V will have no first order sensitivity to errors in the estimation of the phase reference. Since this is an estimation of an unknown parameter, there is no valid *a priori* probability associated with V. The best estimate is found simply by maximizing λ with respect to V to obtain the *Maximum Likelihood* estimator $\hat{V}$. This can be accomplished by setting to zero the derivative of λ in (3) with respect to V.

$$0 = \frac{d\lambda}{dV} = -\frac{1}{2n^2}\sum_k 2(x_k - s(1+V\,cos(2\pi\nu\tau_k)))(-s\,cos(2\pi\nu\tau_k)) \tag{4}$$

2.2. Estimation of V using the *detection phasor*

The solution to the maximization of the likelihood function can be found by forming the following sufficient statistic:

$$d = \sum_k (x_k - \hat{s})\ cos(2\pi\nu\tau_k) \tag{5}$$

where $\hat{s}$ is an estimate we have formed of the actual value of s, the photometric signal strength. Then d will be a direct estimator of V when properly normalized:

$$\hat{V} = \frac{d}{\mathcal{S}} \tag{6}$$

where

$$\mathcal{S} = \hat{s} \sum_k cos^2(2\pi\nu\tau_k) \ . \tag{7}$$

In the case that the phase $2\pi\nu\tau_k$ is totally random (modulo 2π), then for a large summation the normalization approaches an expected value given by the following simple form:

$$\mathcal{S} \approx \frac{K}{2}\hat{s} \tag{8}$$

where K is the number of data used in the summation (5). In this case, errors in the estimation of $\hat{s}$ used in the subtraction of the average signal in (5) will totally cancel, although normalization of the estimator using (7) or (8) will still be affected proportionately to errors in $\hat{s}$.

Estimator error due to measurement noise at the assumed RMS level n can be ascertained by noting that the magnitude of the likelihood has dropped to $1/\sqrt{e}$ of its peak value when the parameter V is at one standard deviation from the solution (8) found for $\hat{V}$. This occurs when $\lambda = \lambda_{peak} - \frac{1}{2}$ corresponding to

$$\frac{1}{2}\frac{d^2\lambda}{dV^2}(V - \hat{V})^2 = -\frac{1}{2} \ . \tag{9}$$

The second derivative of λ with respect to V around the maximum point is easily found to be:

$$\frac{d^2\lambda}{dV^2} = -\frac{s}{n^2}\mathcal{S} \ . \tag{10}$$

Again under the condition of a totally random reference phase, we find the expected RMS error ϵ in the estimation of V due to measurement noise to be:

$$\epsilon = \frac{n}{s}\sqrt{\frac{2}{K}} \ . \tag{11}$$

In the case of V representing a *complex* correlation, it is not hard to see that (5) will stand as an estimator of the *real part* of V and that the imaginary part of V could similarly be estimated as:

$$\hat{V}_i = \frac{1}{\mathcal{S}'}\sum_k (x_k - \hat{s})\ sin(-2\pi\nu\tau_k) \quad \text{where :} \quad \mathcal{S}' = \hat{\mathrm{s}} \sum_{\mathrm{k}} \sin^2(2\pi\nu\tau_{\mathrm{k}}) \ . \tag{12}$$

Again in the case of a totally random phase ψ we find that $\mathcal{S}' \approx \mathcal{S} \approx \frac{K}{2}\hat{s}$ and the estimator for complex V can be written in the more pleasing form:

$$\boldsymbol{\hat{V} = \frac{2}{K\,\hat{s}}\mathcal{D}^*} \tag{13}$$

where we define the complex *detection phasor* $\mathcal{D}$ as:

$$\boldsymbol{\mathcal{D} = \sum_k (x_k - x_0)e^{j2\pi\nu\tau_k}} \ . \tag{14}$$

In this case, assuming a totally random phase $2\pi\nu\tau_k$, an error in x_0, the D.C. term subtracted from each data point, will conveniently cancel in the summation (14). This is a particularly attractive property especially when using detectors whose D.C. offset is large compared to the signal.

The detection phasor $\mathcal{D}$ thus defined forms the basis for the coherent estimators and refinements to be presented in this paper. The mathematics presented can also be shown to encompass as special cases more specific estimators of fringe visibility, as described in Appendix A.

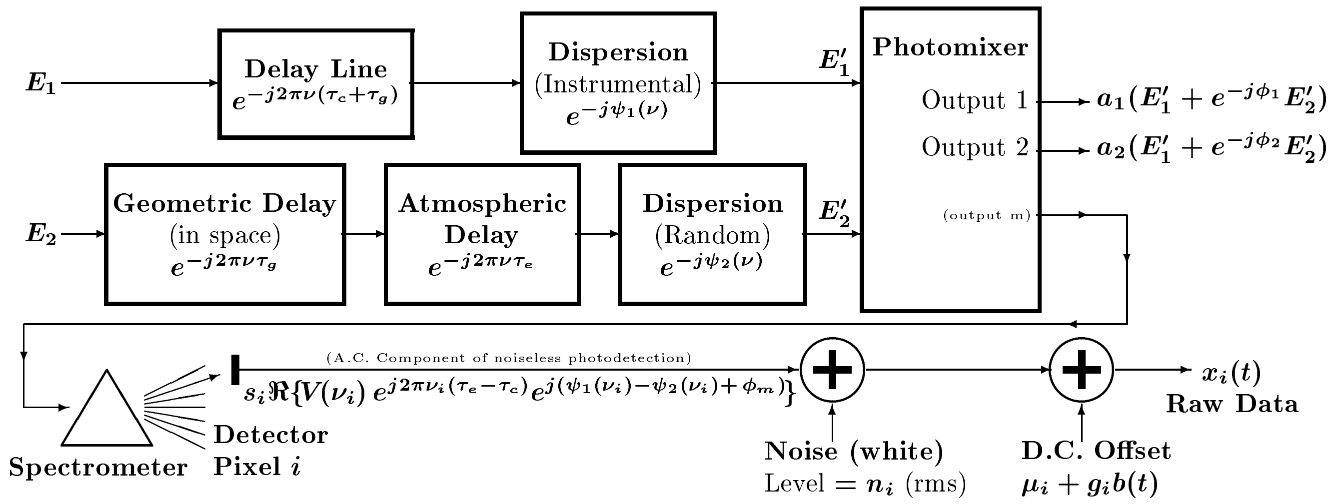

Figure 2. Signal flow model for the production of raw data observed at detector channel i observing light centered around ν_i. Common-mode delays are omitted. Hardware delay line compensates for geometrical delay, and may introduce instrumental delay τ_c. Optical phase is additionally affected by random atmospheric delay τ_e, instrumental dispersion ψ_1 and random components of dispersion ψ_2. Light is mixed in relative phase θ_m at output m of the photomixer to be spectrally dispersed and detected. Gaussian white noise and detector D.C. offsets are additive in the production of the raw data x_i.

3. DETAILED MODEL AND PREPROCESSING

Consider a general two-element single-mode interferometer. Two samples of the stellar wavefront in space pass through different sections of the atmosphere to be received by two telescopes. Both light paths experience a common delay in traversing the atmosphere, but due to atmospheric inhomogeneities the light in path 2 experiences an additional delay τ_e. In our simulations we have modeled τ_e according to the expected $f^{-8/3}$ power spectrum predicted according to Kolmogorov turbulence; the *a priori* statistics of τ_e are useful in optimizing the fringe tracking estimator,[5] but do not directly enter into the estimation of fringe visibility V given an unbiased estimator for τ_e.

Due to its position on the sky with respect to the interferometer's baseline, there is an additional *geometric* delay denoted τ_g. The light from telescope 1 passes through a delay line which will match the geometric delay τ_g and introduces an additional controlled delay τ_c which may be used to (partially) correct for the atmospheric delay τ_e experienced by the light received at E_2, and/or to introduce an intentional phase modulation in the detected interference.

The electric fields are combined in the *photomixer* (or "beam combiner") in two or more phases to produce interference, and the interfered light may be spectrally dispersed prior to detection. The light reaching each detector pixel is recorded as a function of time, either as an analog signal corrupted by gaussian noise, or as a series of photon events corrupted by dark counts (we will assume analog detection for the balance of this paper, but the use of photon-counting detectors leads to a similar analysis[6]). Realtime processing of the interferometric data may (but need not always) be used to estimate the uncorrected component of OPD error in order to servo the delay line setting τ_c.

A more detailed model of the interferometric signal as detected by a single detector channel (or pixel) is depicted in Figure 2. In addition to the geometric and atmospheric delays, we mention here additional phase shifts due to dispersion related to the additional atmospheric delay (which is deterministic given τ_e), dispersion from water vapor fluctuations and pseudo-dispersion due to differential refraction[7,3,4] (both of which are zero-mean random processes), dispersion caused by passage through an air-filled delay line (deterministic given the delay line setting), and other constant instrumental delays and phase shifts which can be calibrated. Finally the electric fields of the two beams are mixed in a particular phase depending on which of the two (or more) outputs of the photomixer is being considered; for a lossless two-output photomixer these are normally expected to be opposite phases (separated by π radians). In the following analysis we shall not address the random components of dispersion affecting the optical path; these can have a considerable effect over a wide band, particularly at visible wavelengths. We will simply lump all of the calibratable phase shifts, including the mixing phase of the photomixer, into a net detection phase for a detector channel (or pixel) labeled i which we will call θ_i; even though we do not write this as a function of time, introducing a known time dependence will not fundamentally change any of our conclusions. So we assume that detector channel i receives light selected by the spectrometer or filter centered at optical frequency ν_i with a symmetric bandpass function given by $G(\nu - \nu_i)$ interfered in the phase θ_i.

Then we can write the response function of detector channel i based on this model (assuming an analog detector).

$$x_i(t) = (\text{DC offset}) + (\text{noise}) + \gamma_{\rm i}(\rm t)\, s_i \mathrm{Re}\left\{ e^{j\theta_i} e^{j2\pi\nu_i\tau(t)}\, V(\nu_i) \right\} \tag{15}$$

Data of this type are recorded at regular time intervals of Δt in each of M detector channels.

In this equation s_i represents the photometric flux of the astronomical object which might be measured, for instance, by chopping against the sky background. $\gamma_i(t)$ is a visibility reduction factor that will be detailed below. These two factors multiply the magnitude of V to yield the net magnitude of interference measured by the detector. Therefore it is clear that a measurement of s_i is required to obtain a normalized visibility $V(\nu_i)$ and that errors in that photometric measurement will generate proportional errors in the visibility estimate. An alternative formulation would simply be to measure the product $s_i V(\nu_i)$ without normalization and to use that result directly in image or parameter reconstruction, as is standard in radio interferometry. All the following mathematics can be interpreted in terms of the unnormalized correlation V simply by setting s=1. In that case the $V(\nu)$ will be expressed in units of (correlated) power, rather than a dimensionless coefficient of "fringe contrast."

The DC offset term present in the above model will have a somewhat different significance depending on the nature of the detector and the dominant source of noise. In the case of an analog detector in the near infrared or visible where the detector dark current is small, this term might mainly reflect the additive effect of the object's photometric strength given by s_i. Then observation of this offset (perhaps by averaging two interferometric channels of opposite mixing phase) may be used, following subtraction of the known dark current, as an estimator for s itself during an observation. At longer wavelengths where we expect thermal background emission to dominate, it would be advantageous to model the offset term for a detector channel as a constant (or slowly varying) term μ_i plus one proportional to background fluctuations $b(t)$ common to all detection channels, scaled by a channel gain g_i found to apply to such fluctuations in channel i. Then we could rewrite (15) for measurement time t_k as:

$$x_{i,k} = x_i(t_k) = \mu_i + g_i b(t_k) + (\text{noise}) + \gamma_{\rm i,k}\, s_i\, \mathrm{Re}\left\{ e^{j\theta_i} e^{j2\pi\nu_i\tau(t)}\, V(\nu_i) \right\} \quad . \tag{16}$$

g_i, the channel gain relative to background fluctuations, can first be determined as an instrumental parameter. Then during an interferometric observation, by averaging the fluctuations common to all detector channels, an estimate of b(t), the varying background level can be obtained, and used to estimate this term in (16) so that it can be subtracted in the pre-processing stage.

3.1. Amplitude calibration of visibility degradations

Due to the multitude of effects which can potentially reduce the measured visibility, the net response of an interferometer must inevitably be calibrated empirically by observing an unresolved star (or one with a known visibility); this is typically done in temporal proximity to and under similar conditions to the scientific observation. However many of the sources of visibility reduction can not only be estimated statistically, but monitored concurrent with the observation, generating a more specific calibration with a higher degree of accuracy. For instance, by monitoring the fluctuating optical coupling coefficient into the fibers whose light gets interfered, the FLUOR detector has been successful in obtaining visibility measurements of unparalleled accuracy.[13] Here we shall describe a number of sources whose effect on the measurement of visibility can be inferred from the data itself and likewise used to generate specific corrections for a given set of data. We shall lump these corrections into a coefficient $\gamma_{i,k}$ which reflects all such reductions in visibility affecting the raw datum $x_{i,k}$ generated at time t_k in detector channel i.

The photometric correction discussed above, compensates for fluctuations in the amount of light accepted by a spatial filter, perhaps monitored by additional detectors as is done in FLUOR (such a photometric monitoring system will also be used optionally in the MIDI instrument, Section 1.2). If P_1 and P_2 are the photometric levels of the interfering beams, then the instantaneous amplitude of interference due to these fluctuations will be degraded by $\gamma^{(P)}$ given by $\sqrt{P_1 P_2}/P_0$ where P_0 is the nominal optical power received in either beam. If these photometric signals are not measured in real time, then the correction can only be applied statistically using models of that fluctuation.[6]

An important source of visibility reduction is encountered when coherently integrating visibilities using a phase reference from an *estimate* $\hat{\tau}$ of the differential OPD τ. In the case of an estimator whose estimation error $\Delta\tau = \hat{\tau} - \tau$ is zero mean gaussian, the visibility reduction $\gamma^{(\Delta\tau)}$ is easily shown to be $exp(-2\pi^2\nu^2\epsilon_\tau^2)$ where ϵ_τ^2 is the mean squared OPD estimation error. If well modeled, that error level can be accurately quantified,[5] and changes in the magnitude of atmospheric turbulence can be monitored and introduced into the model to track changes in that quantity. Even in the case of a probabilistic determination for $\tau(t)$ which is not gaussian, or even multi-modal (consisting of several ambiguous possibilities), it is still possible to use such a phase reference for coherent integration, requiring a more complicated procedure.[6]

One visibility reduction that could be calculated accurately (assuming a continuum object spectrum) results from the deterministic effect of a range of interference phase due to a specific OPD error in conjunction with the finite width of a filter or spectral channel in a spectrally dispersed detection system. In the case of a spectral passband $G(\nu)$ having a boxcar shape of width $\Delta\nu$, this is easily calculated, for an OPD offset τ as $\gamma^{(\Delta\nu)} = sinc(\Delta\nu\ \tau)$. However with any arbitrary passband $G(\nu)$, the particular function for $\gamma^{(\Delta\nu)}$ can be found as an instrumental characteristic and applied as a function of the tracked OPD offset. It should be noted that if there is any measure of differential dispersion in the optical system, then the effective τ that must be used in this calculation must also include the group delay due to that dispersion at that optical frequency.[6] Since any such dispersion must have been calibrated in a multiwavelength coherent detection system, this instrumental characteristic will already be known.

Another visibility reduction that can be well modeled and applied is due to the *slope* of the differential OPD which induces a change in the phase of interference during the exposure of a detector element before it is read out. Again this can be continuously modeled using the phase reference. A slope in the OPD of $\frac{d\tau}{dt}$ during an exposure of duration Δt at optical frequency ν will reduce the visibility by $\gamma^{(\Delta t)} = sinc(\nu \frac{d\tau}{dt} \Delta t)$. Although a more detailed correction for the higher derivatives of τ could in principle be applied, these will be insignificant for short exposures.

Finally, any reductions in the magnitude of fringe visibility not directly accounted in such a measurable way, will simply be lumped into a final empirical calibration $\gamma^{(cal)}$. This is the quantity that would be adjusted to yield the correct result when a "calibrator" source is observed in conjunction with a scientific observation. Then for each data point at time t_k in detector channel i, we can estimate a net visibility degradation factor $\gamma_{i,k}$ based on these (and other) factors. Using the specific forms calculated above, for instance, we would find the net coefficient to be:

$$\gamma_{i,k} = \gamma^{(P)}(t_k) \ \cdot \ \gamma^{(\Delta\tau)}(\nu_i, \epsilon_\tau^2(t_k)) \ \cdot \ \gamma^{(\Delta\nu)}(\Delta\nu_i, \tau(t_k)) \ \cdot \ \gamma^{(\Delta t)}(\nu_i, \tfrac{d\tau}{dt}) \ \cdot \ \gamma^{(cal)} \tag{17}$$

$$= \frac{\sqrt{P_1(t_k)P_2(t_k)}}{P_0} \ e^{-2\pi^2\nu_i^2\epsilon_\tau^2(t_k)} \ sinc(\Delta\nu_i \ \tau(t_k)) \ sinc\Big(\nu_i \frac{d\tau(t_k)}{dt}\Delta t\Big) \ \gamma^{(cal)} \quad . \tag{18}$$

3.2. Preprocessing

Let us first transform the raw data obeying the model described above (16) in order to create preprocessed data in which the offset biases have been removed and the data optimally weighted for subsequent processing. The estimate of the varying background radiation's contribution to the reading in detector channel i, given by $g_i\hat{b}(t_k)$, is subtracted from the raw data An estimate of the offset applying to that particular detector channel, μ_i, is also subtracted. After removal of the estimated offset, the result is divided by n_i, the RMS white noise level in that channel to produce data having a noise level of unity. Then each datum is weighted by its own signal to noise ratio $\gamma_{i,k}s_i/n_i$. A final weighting factor $w_{i,k}$ may also be applied. We shall assume initially that all $w_{i,k} = 1$, but if we choose to apply unequal weighting to the data we will stipulate that all $w_{i,k} \leq 1$ so that the "strength" $\mathcal{S}$ defined below will not overestimate the quality of the data. The preprocessed data stream $z_{i,k}$ is thus produced from the raw data $x_{i,k}$.

$$\boldsymbol{z_{i,k} = (x_{i,k} - \hat{\mu}_i - g_i\hat{b}_i)\gamma_{i,k} \frac{s_i}{n_i^2} w_{i,k}} \tag{19}$$

Each preprocessed datum will then have an expected value given by

$$E\{z_{i,k}\} = \gamma_{i,k}^2 \frac{s_i^2}{n_i^2} w_{i,k} \ Re\left\{e^{j\theta_i} e^{j2\pi\nu_i\tau(t)} \ V(\nu_i)\right\} \tag{20}$$

and an RMS noise level which we will assume to be $\gamma_{i,k}\ s_i/n_i$ (although it may be lower if we have chosen $w_{i,k} < 1$).

Using some subset of the preprocessed data $z_{i,k}$ we can again form the detection phasor $\mathcal{D}$ as:

$$\mathcal{D} = \sum_{i,k} e^{j\theta_i} e^{j2\pi\nu_i\tau(t_k)} \; z_{i,k} \quad . \tag{21}$$

Again assuming a totally random net phase $\theta_i + 2\pi\nu_i\tau(t_k)$, we can write the maximum likelihood estimator for complex V in terms of the detection phasor $\mathcal{D}$ as:

$$\hat{V} = \frac{2}{\mathcal{S}} \mathcal{D}^* \quad . \tag{22}$$

We refer to the normalization $\mathcal{S}$ as the *strength* of the estimate; it is given by:

$$\mathcal{S} = \sum_{i,k} \gamma_{i,k}^2 \frac{s_i^2}{n_i^2} w_{i,k} \quad . \tag{23}$$

4. ERRORS RESULTING FROM ASYMMETRY IN SAMPLING PHASE

A required condition mentioned above in order for the estimator (22) to be valid, is that the net interference phase $\psi_{i,k}$ for the set of data used in producing a particular estimate must be uniformly distributed over $[0, 2\pi]$. This is the combined phase of the instrumental phase for detector channels θ_i and the effect of the OPD given by $2\pi\nu\tau(t_k)$. Actually somewhat weaker conditions on the phase statistics are required for this purpose.

Let us define the N^{th} order net phase statistic P_N applying to the detection phasor formed from a subset of preprocessed data $z_{i,k}$ over k (time) and also i (detector channel) as:

$$P_N = \sum_{i,k} e^{j\,N\,\psi_{i,k}} \; \gamma_{i,k}^2 \frac{s_i^2}{n_i^2} w_{i,k} \quad \text{where:} \quad \psi_{\mathrm{i,k}} = \theta_{\mathrm{i}} + 2\pi\nu_{\mathrm{i}}\tau(\mathrm{t_k}) \quad . \tag{24}$$

Note that the 0^{th} order phase statistic P_0 is identical to the estimator strength $\mathcal{S}$ defined above. For the purposes of phase symmetry we would like to null the first and second order phase statistics:

$$P_1 = \sum_{i,k} e^{j\,\theta_i} e^{j\,2\pi\nu_i\tau(t_k)} \gamma_{i,k}^2 \frac{s_i^2}{n_i^2} w_{i,k} \quad ; \quad P_2 = \sum_{i,k} e^{j\,2\,\theta_i} e^{j\,4\pi\nu_i\tau(t_k)} \gamma_{i,k}^2 \frac{s_i^2}{n_i^2} w_{i,k} \quad . \tag{25}$$

There are three ways in which we may be able to force these summations to zero. First, the design of the photomixer (beam combiner) can be tailored for this purpose by creating symmetry in the phases when integrated over a set of detector channels i. Second, one can introduce into the instrumental delay τ_c a modulation term designed to cause phase symmetry when integrated over detection time bins k. Third, if we start out with a data set almost possessing phase symmetry, we can make small adjustments to the supplementary weighting factors $w_{i,k}$ in order to null P_1 and/or P_2 in (25).

Although not required, it will clarify the discussion if we require the weighting factors $w_{i,k}$ to be separable with respect to detector channel i, and time bin k: $w_{i,k} = w'_i \; w''_k$. We would also like to write the visibility amplitude calibration factor γ to only be a function of time, which will generally be true when the detector channels are similar and at about the same optical frequency, thus: $\gamma_{i,k} = \gamma_k$. We can then write the N^{th} order phase statistic as a product of a sum over detector channels i, and a sum over time bins k.

$$P_N = \left[\sum_i e^{jN\theta_i} \frac{s_i^2}{n_i^2} w'_i \right] \cdot \left[\sum_k e^{jN\,2\pi\nu_i\tau(t_k)} w''_k \, \gamma_k^2 \right] \tag{26}$$

We note that the first summation will be dependent on the photomixer design, and can be adjusted by varying the channel weighting w'_i. The second summation can be reduced by introducing a suitable delay modulation scheme, and could be further influenced by adjustment of the weighting factors w''_k.

4.1. First order phase asymmetry and sensitivity to background fluctuations

If we assume that all of the outputs of a photomixer, indexed by i, are employed equally in forming the estimator, then, in the case of a lossless photomixer, we are assured mathematically that P_1 will be zero.

Of course in practice we cannot expect that cancellation to be complete. A first order phase asymmetry will have the effect of causing the detection phasor (21) to be sensitive to small variations in DC offset terms common to all detector channels. In the case of a mid-infrared instrument, the background radiation level b can be thousands of times stronger than the signal, so even small fluctuations in b could generate large estimation errors.

To quantify the effect, let us introduce our model for the data received at detector channel i at time t_k given by (16) into the expression for the detection phasor (21) under the condition of no actual interference (equivalent to setting V=0). Let us take the weighting factors to be in the above separable form and stipulate that we are creating an average using only detector channels receiving light of the same wavelength. No subtraction for channel offset μ_i or background radiation $g_i b(t)$ has been attempted. Then we shall find the expected value the detection phasor in order to illustrate its sensitivity to these offsets.

$$E\{\mathcal{D}\} = \sum_i \sum_k e^{j\theta_i} e^{j2\pi\nu\tau(t_k)}\ \gamma_k \frac{s_i}{n_i^2} w_{i,k}\ (g_i b(t_k) + \mu_i)$$

$$= \left[\sum_i e^{j\theta_i} \frac{s_i}{n_i^2} w_i'\ g_i\right] \cdot \left[\sum_k e^{j2\pi\nu\tau(t_k)} \gamma_k w_k''\ b(t_k)\right] + \left[\sum_i e^{j\theta_i} \frac{s_i}{n_i^2} w_i'\ \mu_i\right] \cdot \left[\sum_k e^{j2\pi\nu\tau(t_k)}\ \gamma_k\, w_k''\right] \qquad (27)$$

Since the variable background level $b(t)$ is contained only in the second factor of the first term, its effect on $\mathcal{D}$ can be cancelled by forcing the sum over detector channels i in the first factor to be zero. This summation includes g_i, the response coefficient of a detector channel to the variable background radiation $b(t)$. However in the expected case that the detector channels involved are looking at the same optical signals combined in a different phase, then the responsivity of the detector to background radiation given by g_i and the response of that channel to the signal given by s_i will stay in proportion. Then nulling of this term can be seen to be equivalent to nulling P_1 by forcing the first summation in (26) to zero. This can generally be accomplished by slightly adjusting the weighting factors w_i'.

On the other hand, sensitivity of $\mathcal{D}$ to unknown variations in μ_i, the offset for a single detector element, requires nulling of the temporal aspect of the first order phase asymmetry contained in the second factor in brackets in the second term of (27). Since the OPD $\tau(t)$ contains both an atmospheric term and an instrumental term, it is possible to control the instrumental delay in such a way so as to randomize the net phase term. For instance, the so-called ABCD method (Appendix A) repeatedly modulates the instrumental delay in four $\pi/2$ phase steps which tends to cause this summation to approximate zero. However that will only hold exactly true at one optical frequency so that a specific delay modulation optimized for that wavelength will be suboptimum at different wavelengths. What's more, even though the sum will have been forced to zero given, say, a constant value of τ, any *specific* sum based on a particular time series $\tau(t)$ at a specific optical frequency will not be perfectly nulled. However it would generally be possible to make fine adjustments to the temporal weighting coefficients w_k'' so that the sum could be forced exactly to zero, a different set of weighting coefficients being necessary for estimations of V at different optical frequencies. The choice of weighting coefficients for a long sum to null this factor will not be unique, and no specific method is presently proposed to accomplish this in an optimum manner.

As an example, the atmospheric OPD shown in Figure 3 upper left has been used to create simulated data in 96 detector channels corresponding to an unresolved star: $V = 1$, $arg(V) = 0$ for all ν. To magnify the problem, only one phase of interference has been generated. With perfect background subtraction a good estimate of V is produced in Figure 3 upper right. However if an uncorrected D.C. offset of 5 times the RMS noise level is introduced, the same estimator exhibits wide fluctuations from the underlying V as seen in Figure 3 lower left. If we introduce "ABCD" modulation, so that the instrumental OPD is stepped one quarter of a wavelength every measurement, repeated after every four detections, then this source of instrumental error has been greatly reduced as seen in Figure 3 lower right. Note that the correction is superb at the optical frequency of 30 Terahertz which corresponds to the 10 micron wavelength on the basis of which the "ABCD" steps were generated, whereas substantial inaccuracies remain at widely different wavelengths.

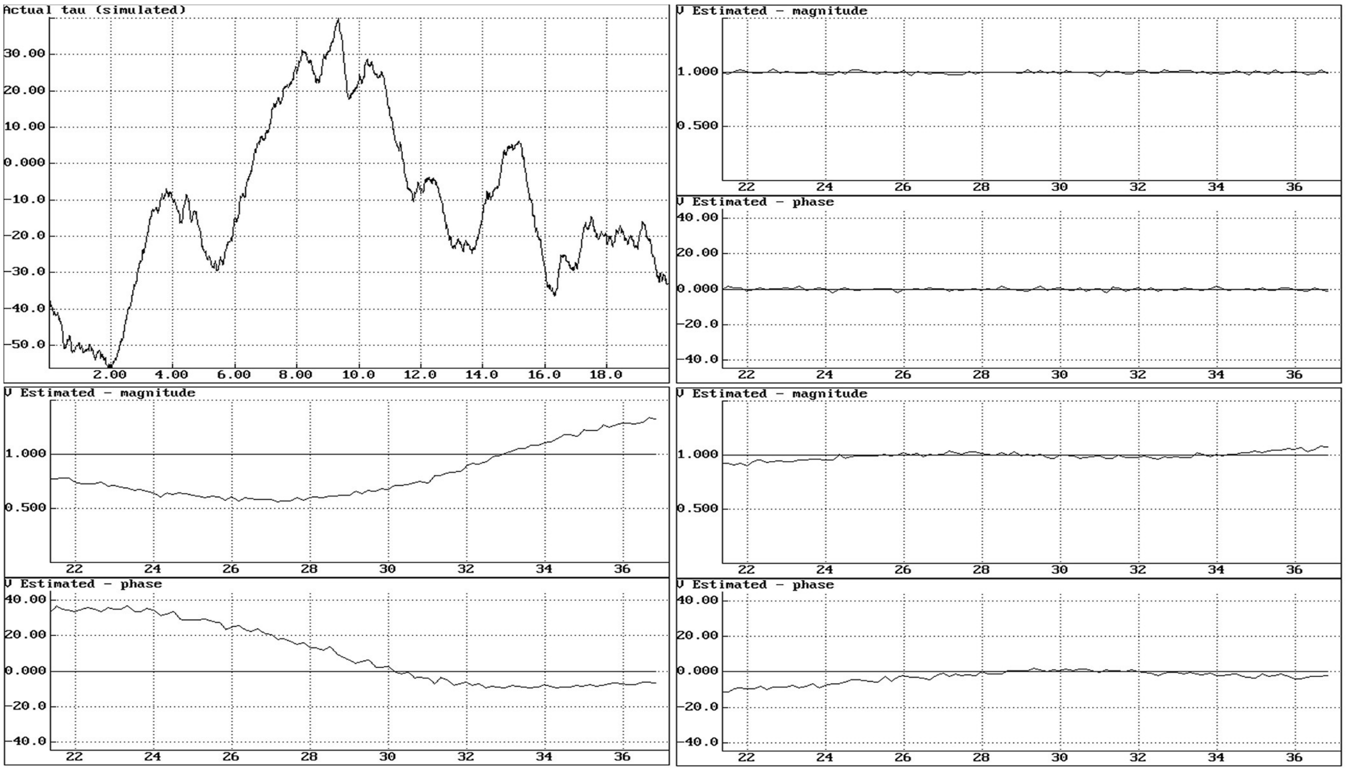

Figure 3. An atmospheric OPD function (upper left) plotted in femtoseconds over a 20 second period is used to simulate interferometric data from which a good estimation of visibility is made (upper right). However the addition of an uncorrected D.C. offset in the data creates wild errors in the estimation due to first order phase asymmetry (lower left) which is largely cancelled when an intentional "ABCD" modulation is added to the OPD (lower right) especially around 30 Terahertz corresponding to the 10 μm wavelength for which the OPD steps were exactly $\lambda/4$.

4.2. Second order phase asymmetry

Unlike first order phase asymmetry which we have seen will cause an estimator using the detection phasor $\mathcal{D}$ to be sensitive to DC offsets in the detector outputs, non-zero values of P_2 will cause the estimator to have unequal, and incorrect sensitivities to the real and imaginary parts of the underlying V^*. For instance, 20 seconds of interferometric data from a two-phase detection system with 96 detector channels (48 spectral channels) has been simulated and used to estimate $V(\nu)$ using (22); the underlying model again is based on $V = 1$, $arg(V) = 0$. The estimate shown in Figure 4, upper left, (along with the assumed atmospheric OPD) is seen to have very substantial errors due to the large second order phase asymmetry in this short observation.

It is again possible to null P_2 using the methods discussed above to null P_1 by varying the weighting coefficients and/or tailoring the instrumental OPD to null the second sum for P_2 in (26) (again the "ABCD" OPD modulation scheme accomplishes this in a statistical sense). For instance, by introducing "AB" modulation, in which each alternate exposure has an added instrumental OPD of zero or $\lambda/4$, the errors seen in the upper left of Figure 4 have been reduced almost by a factor of 10 as shown in the upper right graphs (note the change in the vertical scale). The phase reference used, in which the 8 femtosecond "AB" modulation has been added to the atmospheric OPD, is shown as well. A similar error reduction is also seen in the lower left of Figure 4, where the integration period has been extended from 20 to 2000 seconds to provide an averaging of phase statistics which would be expected to reduce the effect of P_2 by a factor of 10.

However extending the observation of a strong signal by this amount should never be necessary. In fact, suppression of P_2 may be unnecessary altogether inasmuch as its effect is predictable, unlike the case for non-zero P_1 which made the estimator sensitive to background offsets and fluctuations which may not be directly observable

*More precisely, this will apply to the real and imaginary parts of $Ve^{j\psi}$ where $\psi = \frac{1}{2}arg(P_2)$.

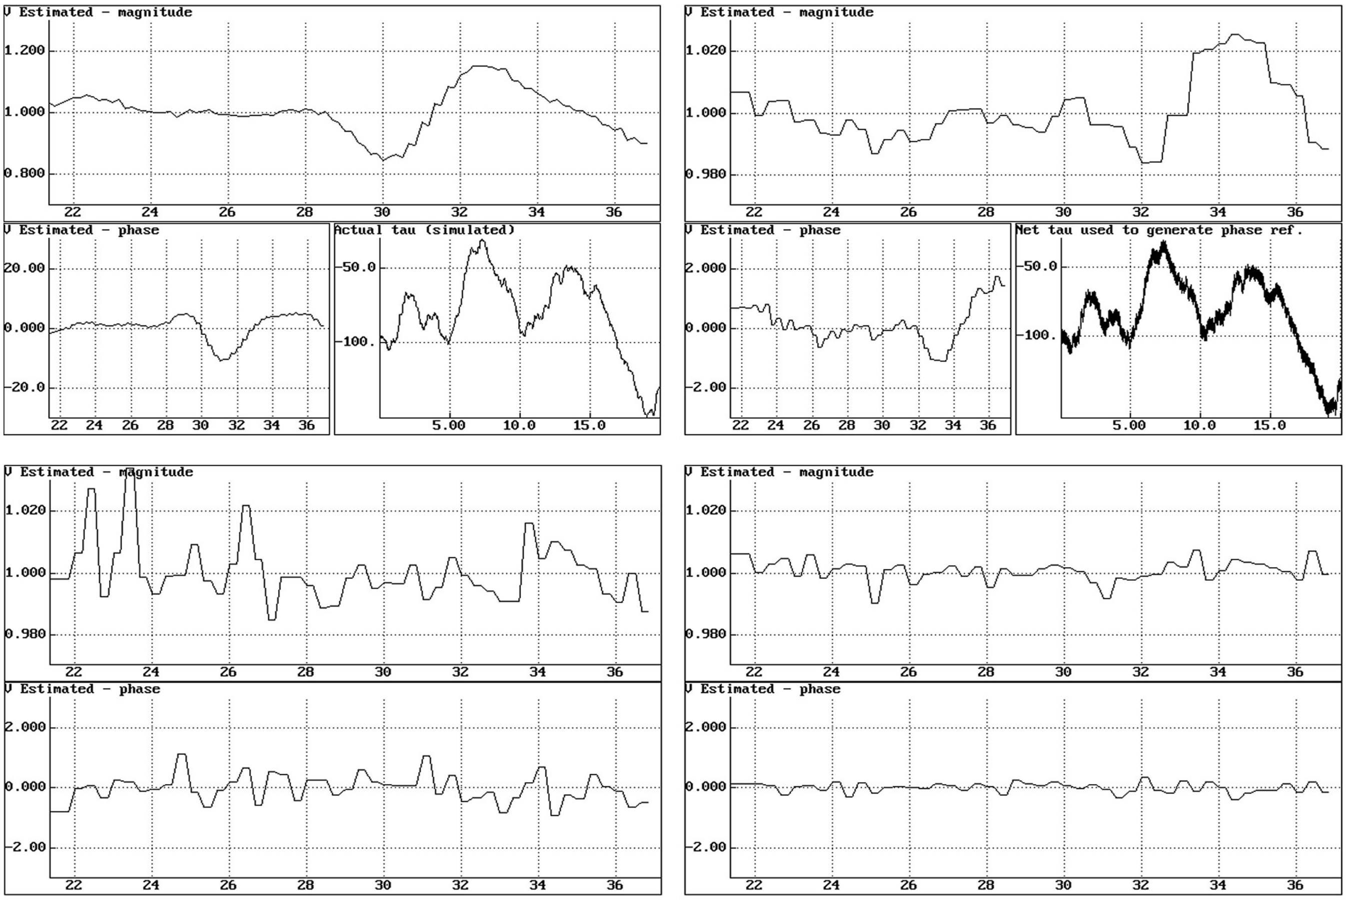

Figure 4. Large errors in the estimation of $V(\nu)$ using (22) (upper left) are greatly reduced (note change in vertical scale) by using $\lambda/4$ square wave OPD modulation (upper right), or a 2000 second observation (lower left). Best results are obtained (lower right) by applying the estimator (29) which corrects for the effect of non-zero P_2.

simultaneously with an observation. It can be shown that, knowing P_2, the complex detection phasor will have an expected value given by

$$E\{\mathcal{D}\} = \frac{1}{2}V^*\mathcal{S} + \frac{1}{2}VP_2 \quad . \tag{28}$$

We can then produce a modified estimator for V based on the detection phasor $\mathcal{D}$:

$$\hat{V}(\nu) = 2\,\frac{\mathcal{S}\,\mathcal{D}^* \;-\; P_2^*\,\mathcal{D}}{\mathcal{S}^2 \;-\; |P_2|^2} \quad . \tag{29}$$

This estimator now includes a correction for second order phase asymmetry given by P_2 thereby canceling its effect.

In the above simulation, the errors seen in the upper left of Figure 4 have been totally suppressed in the lower right graphs using the estimator (29) incorporating the correction from P_2. The error reduction accomplished using this formula appears to be complete and was not augmented by additional application of the "AB" modulation or a long integration according to further simulations (not shown).

5. ESTIMATION ERROR

We earlier computed (11) the expected error in estimating V based on probability theory and a simplified model. Applying the detailed model (15), preprocessing (19), and definition of the strength $\mathcal{S}$ (23), it can be shown that the error level due to measurement noise in the estimator is determined solely by $\mathcal{S}$, the estimator's "strength" (hence the name). In the case of second order phase symmetry, the RMS errors ϵ_r and ϵ_i in the estimations of the visibility's real and imaginary parts respectively are given by

$$\epsilon_r = \epsilon_i = \sqrt{\frac{2}{\mathcal{S}}} \quad . \tag{30}$$

In the more general case, if $\psi = \frac{1}{2}arg(P_2)$ then the errors in the real and imaginary parts of $\hat{V}e^{j\psi}$ will satisfy:

$$\frac{1}{\epsilon_r^2} + \frac{1}{\epsilon_i^2} = \mathcal{S} \qquad . \tag{31}$$

The strength $\mathcal{S}$ can be improved, and the error lowered, by increasing the duration of an observation or by averaging over a greater number of detector channels. For instance, the error level of .013 RMS seen in the simulation results of Figure 1 could have been reduced to .005 by increasing the observation time from 20 to 135 seconds. The same decrease in the error of measuring $V(\nu_1)$ could be achieved by averaging over 7 spectral channels around ν_1. That averaging is optimally performed by extending the range of i in the formation of the detection phasor (21) so that all contributions are weighted by their own signal to noise ratio. Of course this is at the expense of resolution in the optical frequency domain, however in many cases we would expect the visibility curve to be a smooth function of wavelength, and such averaging would be acceptable.

6. ESTIMATION OF V IN THE CASE OF $1/F$ DETECTOR NOISE

The estimator we have derived is optimized for the case of the noise term in each detector channel in (15) having white spectral characteristics. In other words the noise terms affecting exposures at different times are uncorrelated. This may be the result of shot noise either from the signal photons, background photons, or dark current. In addition to a white noise term, however, real detectors may suffer from low-frequency dominated noise (or so-called *1/f noise*) which will not be optimally rejected by the previous estimator.

It may be possible to tailor estimators of V given an expected noise spectrum which is not white. We will now develop an estimator which is optimized for noise whose power spectrum follows a $1/f^2$ power law (the Weiner-Levy process, or "random walk"). We will find that the resulting estimator is itself not far from optimum even when used in the presence of white noise, suggesting that its performance may not be far from optimum when used in a system whose detectors generate low-frequency dominated noise following a $f^{-\alpha}$ power law for other α in the range $0 < \alpha < 2$, though this presumption has not been rigorously analyzed.

The approach to be followed assuming detector noise given by a $1/f^2$ power law, would therefore model the *difference* of noise terms present in subsequent readouts of a detector channel i as being independent gaussian random variables with a variance of n_i^2. In this case we will alter the pre-processing of the raw data, previously given by (19). The new intermediate data $z_{i,k}$ will be computed by forming the *difference* between subsequent raw data in a detector channel, again divided by the RMS noise level and weighted by the signal to noise ratio and a supplementary weighting factor $w_{i,k}$.

$$z_{i,k} = (x_{i,k} - x_{i,k-1})\, \gamma_{i,k} \frac{s_i}{n_i^2} w_{i,k} \tag{32}$$

This produces a sequence of data to operate on whose noise spectrum is again white. We will find in this case that while the noise in the $z_{i,k}$ is governed by the detector characteristics, the strength of the preprocessed signal $z_{i,k}$ is dependent on modulation of the net OPD $\tau(t)$, and has an amplitude proportional to the magnitude of $q_{i,k}$ defined as:

$$q_{i,k} = e^{j\theta_i} \left(e^{j2\pi\nu_i \tau(t_k)} - e^{j2\pi\nu_i \tau(t_{k-1})} \right) \qquad . \tag{33}$$

Unlike in the previous case, we will now find that the performance of the estimator is highly dependent on a delay modulation scheme (control of the instrumental OPD τ_c) in which the difference between τ during subsequent optical integrations is a large fraction of a wavelength. In fact the magnitude of q in (33) is maximized when the change in τ between subsequent detector readouts is half a wavelength, but of course that could only be exactly true for a single wavelength. But we can see that even for an interferometer operating over a 2:1 wavelength range, a delay modulation could be performed which would translate to a phase shift between 120 degrees and 240 degrees over that wavelength range (not including the random atmospheric phase shift over Δt). At those extremes, $|q|$ would still attain 86% of its maximum possible value of 2.

In this case q, the phasor difference between the phases of subsequent data, is used instead of the phase itself in the definition of the detection phasor $\mathcal{D}$, now defined as:

$$\boldsymbol{\mathcal{D}} = \sum_{i,k} \boldsymbol{q_{i,k}\ z_{i,k}} \quad . \tag{34}$$

As we did with the basic estimator (29), we will define statistics similar to $\mathcal{S}$ (23) and P_2 (25), which are not dependent on the measured data but only on the statistics of the differential OPD during the period over which those data were taken.

$$\mathcal{S} = \sum_{i,k} |q_{i,k}|^2\, \gamma_{i,k}^2\, \frac{s_i^2}{n_i^2} w_{i,k} \qquad Q_2 = \sum_{i,k} q_{i,k}^2\, \gamma_{i,k}^2\, \frac{s_i^2}{n_i^2} w_{i,k} \tag{35}$$

Again $\mathcal{S}$ is a measure of the "strength" of the entire estimate and is similarly related to the error level achieved. Q_2 again is a measure of second order phase asymmetry. Note that first order phase asymmetry is of no concern following preprocessing using (32) since detector offsets and background fluctuations are naturally removed (as well as can be) when the differences between pairs of subsequent data are formed.

We then arrive at an estimator of the same form as previously derived in (29).

$$\boldsymbol{\hat{V}(\nu) = 2\, \frac{\mathcal{S}\, \mathcal{D}^* \; - \; Q_2^*\, \mathcal{D}}{\mathcal{S}^2 \; - \; |Q_2|^2}} \tag{36}$$

Although optimized for use with an f^{-2} noise spectrum, it appears that even in the original case of raw data corrupted by white noise, this estimator would match the performance of (21) as long as the OPD modulation introduced is close to half a wavelength, in order to maximize q as discussed above.

7. CONCLUSION

Coherent estimation of fringe visibility in an optical/IR astronomical interferometer, is obtained using a phase reference applied either using hardware delay tracking, or offline following the observation. Such methods can increase the sensitivity of a stellar interferometer, and provide the phase of visibility with respect to the phase reference. The estimator derived (29) is based on the the detection phasor $\mathcal{D}$ (21) which uses a delay estimator $\hat{\tau}$ to rotate the phase of the preprocessed data stream $z_{i,k}$ (19). With proper amplitude and phase calibration, accurate determinations of visibility amplitude and phase can be obtained, as shown in simulations (Figure 1).

APPENDIX A. ABCD AND RELATED ESTIMATORS

A popular visibility estimator is based on measurement of interference in four quadrature phases produced by rapid delay stepping. Often referred to as the "ABCD" method, such an estimator can be seen to be a special case of the basic coherent estimator (13) formed above using the detection phasor. A, B, C, and D are the signals measured as the phase is stepped from its initial phase which we will call ψ_0, to $\psi_0 + \pi/2$, $\psi_0 + \pi$, and $\psi_0 + 3\pi/2$. This characterization of the phase delays is justified under the condition that the phase is stepped so quickly that the random atmospheric OPD can be considered to be "frozen" during the period of data acquisition. The initial phase ψ_0 may or may not be known with respect to an external source.

Forming the detection phasor for these four data points we find:

$$\mathcal{D} = \sum_{k=0}^{3} x_k e^{j\psi_k} = Ae^{j\psi_0} + Be^{j(\psi_0+\pi/2)} + Ce^{j(\psi_0+\pi)} + De^{j(\psi_0+3\pi/2)} = e^{j\psi_0}\left[(A-C) + j(B-D)\right] \quad . \tag{37}$$

In the case that ψ_0 is known, perhaps using an external fringe tracker, then we have produced a coherent estimator of V:

$$\hat{V} = \frac{2}{4\,\hat{s}}\, \mathcal{D}^* = \frac{e^{-j\psi_0}}{2\,\hat{s}}\left[(A-C) - j(B-D)\right] \quad . \tag{38}$$

With the coherent estimator in (38) we can repeat the procedure M times to produce M estimates of the complex V and average these complex values, a procedure which is clearly identical to directly forming an estimate from a detection phasor (14) employing all 4M data points.

On the other hand, even if we have no knowledge of the initial phase ψ_0, we can still produce an *incoherent estimator* of $|V|^2$ from these 4 data points by computing the magnitude of (38) and averaging those real numbers:

$$W = \frac{1}{M}\sum_m \left|\hat{V}_m\right|^2 \quad . \tag{39}$$

For very high raw signal to noise ratios, W thus formed as an incoherent average will be a good estimate of $|V|^2$. However, as is well known, the process of forming the magnitude of the measured complex correlation, which involves squaring the real and imaginary parts of $\hat{V}_m$, inevitably introduces a term proportional to the power of the measurement noise present in the raw data A,B,C, and D. At raw signal to noise ratios less than one this bias term is actually dominant, and in any case, must be carefully subtracted from W to produce a useful estimate of $|V|^2$. Thus by averaging the data incoherently, we have now produced an estimator that has a strong first order sensitivity to the noise characteristics of the detector and fluctuations in that noise level, thereby contributing a further source of error.[1] In contrast, the effect of detection noise in the coherent estimator (13) is zero mean (in both its real and imaginary parts) and will be reduced in magnitude by $\sqrt{M}$ when the data set grows by the factor M. This is true of the ABCD estimator (38) when a phase reference for ψ_0 is available to enable coherent averaging.

A similar estimator which can be used coherently or incoherently is the "ABC" method in which three measurements, A, B, and C, are made each following stepping of the instrumental delay by $2\pi/3$. In this case, we likewise find, assuming an initial absolute phase of ψ_0, a coherent estimate based on the detection phasor (14) .

$$\hat{V} = \frac{2}{3\,\hat{s}}\,\mathcal{D}^* = \frac{e^{-j\psi_0}}{3\,\hat{s}}\left[2A - (1+j\sqrt{3})B - (1-j\sqrt{3})C\right] \tag{40}$$

The ABC estimator shares all of the features of the ABCD method, and has a slight advantage with respect to the overhead involved in delay stepping and detector readout. However, it does not appear to have been employed in stellar interferometry.

REFERENCES

1. M. M. Colavita, "Fringe visibility estimators for the Palomar Testbed Interferometer," *Publications of the Astronomical Society of the Pacific*, **111**, pp. 111-117, 1999
2. D. Mozurkewich, K. J. Johnston, R. S. Simon, D. J. Hutter, M. M. Colavita, M. Shao, J. L. Hershey, J. A. Hughes, and G. H. Kaplan, "Phase referenced averaging as a method for decreasing the variance of visibility measurements," *NOAO-ESO Conference on High Resolution Imaging by Interferometry*, European Southern Observatory, Garching bei Munchen, Germany, pp 851-854, 1988
3. M. M. Colavita, "Phase referencing for stellar interferometry," *ESO Conference on High-Resolution Imaging by Interferometry II*, European Southern Observatory, Garching bei Munchen, Germany, pp. 845-851, 1992
4. A. Quirrenbach, D. Mozurkewich, D.F. Buscher, C.A. Hummel, and J.T. Armstrong, "Phase-referenced visibility averaging in optical long-baseline interferometry," *Astronomy and Astrophysics* **286**, pp. 1019-1027, 1994
5. J. Meisner, "Atmospheric delay tracking in a long-baseline optical stellar interferometer," *Optical Engineering*, **35**, pp. 1927-1935, 1996
6. J. Meisner, "Coherent integration of fringe visibility employing probabilistic determinations of atmospheric delay," in *Astronomical Interferometry*, R. D. Reasenberg ed., *Proc. SPIE* **3350** pp. 294-308, 1998
7. M. M. Colavita, M. Shao, and D. H. Staelin, "Two-color method for optical astrometry: theory and preliminary measurements with the Mark III stellar interferometer," *Applied Optics*, **26**, pp. 4113-4122, 1987
8. M. M. Colavita, et. al., "The Palomar Testbed Interferometer," *Astrophysical Journal*, **510**, pp. 505-521, 1999
9. A. Quirrenbach, et. al., "PRIMA: study for a dual-beam instrument for the VLT Interferometer," in *Astronomical Interferometry*, R. D. Reasenberg ed., *Proc. SPIE* **3350**, pp. 807-817, 1998
10. M. Shao, and M. M. Colavita, "Potential of long-baseline infrared interferometry for narrow-angle astrometry," *Astronomy and Astrophysics*, **262**, pp. 353-358, 1992
11. J. T. Armstrong, D. Mozurkewich, T. A. Pauls, and A. R. Hajian, "Bootstraping the NPOI: keeping long baselines in phase by tracking fringes on short baselines," in *Astronomical Interferometry*, R. D. Reasenberg ed., *Proc. SPIE* **3350**, pp. 461-466, 1998
12. C. Leinert, et. al., "10-μ interferometry on the VLTI with the MIDI instrument: a preview," *These proceedings*
13. V. Coudé du Foresto, S. Ridgway and J.-M. Mariotti, "Deriving object visibilities from interferograms obtained with a fiber stellar interferometer," *Astronomy and Astrophysics Supplement Series*, **121**, pp 379-392, 1997