

Sensitive visible interferometry with PAVO

Michael J. Ireland^{*a,b}, Antoine Mérand^{c,d}, Theo A. ten Brummelaar^c, Peter G. Tuthill^b,
Gail H. Schaefer^c, Nils H. Turner^c, Judit Sturmann^c, Laszlo Sturmann^c,
Harold A. McAlister^c

^aDivision of Geological and Planetary Sciences, California Institute of Technology, Pasadena,
CA 91125, USA;

^bSchool of Physics, University of Sydney, NSW 2006, Australia;

^cCenter for High Resolution Astronomy, Georgia State University, PO Box 3965, Atlanta,
Georgia 30302-3965, USA;

^dEuropean Southern Observatory, Alonso de Cordova 3107, Vitacura Santiago, Chile

ABSTRACT

The Precision Astronomical Visible Observations (PAVO) beam combiner is a new concept in visible beam combination, recently commissioned at the CHARA array. By creating spatially-modulated fringes in a pupil plane and then dispersing with an integral field unit, PAVO utilizes the full multi- r_0 aperture of the CHARA array over a standard 50% (630-950nm) bandwidth. In addition, minimal optimized spatial filtering ensures calibration that is in principle as good as using single-mode fibers. We describe the design of and initial results from the PAVO instrument.

Keywords: instrumentation: interferometers - techniques: interferometric - optical interferometry

1. INTRODUCTION

The Precision Astronomical Visible Observations (PAVO) beam combiners at both the Center for High Angular Resolution Astronomy (CHARA) array and the Sydney University Stellar Interferometer (SUSI) are designed to be sensitivity-optimized spectrally-dispersed beam combiners with spatial filtering. The implementation at CHARA is a 3-beam combiner and the most developed implementation, so will be focused-on here.

Prior to the development of PAVO, the state of the art in beam-combination in the visible was either pupil-plane fringe-scanning beam-combiners with a $\sim 10\%$ bandwidth (e.g. Ref 1), dispersed group-delay tracking beam combiners² or dispersed image-plane combination.³ The only design in this set that could utilize apertures $\gg r_0$ was the VEGA/REGAIN design,³ which was not sensitivity optimized due to both the losses at the slit and the relatively low quantum-efficiency of the photon-counting intensified CCD detectors.

The PAVO beam-combiner design aims to take advantage of the recent technological advance of Electron-Multiplied CCD detectors (EMCCDs) with their photon-counting ability and high quantum efficiency, to adequately sample fringes from a multi- r_0 aperture and over a very broad bandwidth. Section 2 outlines the optical design of PAVO and the general concept, Section 3 motivates the use of the visible wavelength regime for sensitivity-optimized beam-combination, Section 4 outlines the process of minimally-optimized spatial filtering, Section 5 outlines the data analysis pipeline, Section 6 presents commissioning data and Section 7 concludes and describes the current and future developments for PAVO.

2. CONCEPT AND OPTICAL DESIGN

The general optical concept for PAVO is illustrated in Figure 1. Refractive optics bring the beams to a common image-plane, using sharp-edged optics to place the beams side-by-side. Spatially-modulated fringes are then formed on a cylindrical lenslet array. The optics from the lenslet array to the (EMCCD) detector have a similar effect to an integral field unit: re-imaging the images formed by the lenslets split into a variety of different wavelengths and in orthogonal polarization states.

*E-mail: M.Ireland@physics.usyd.edu.au

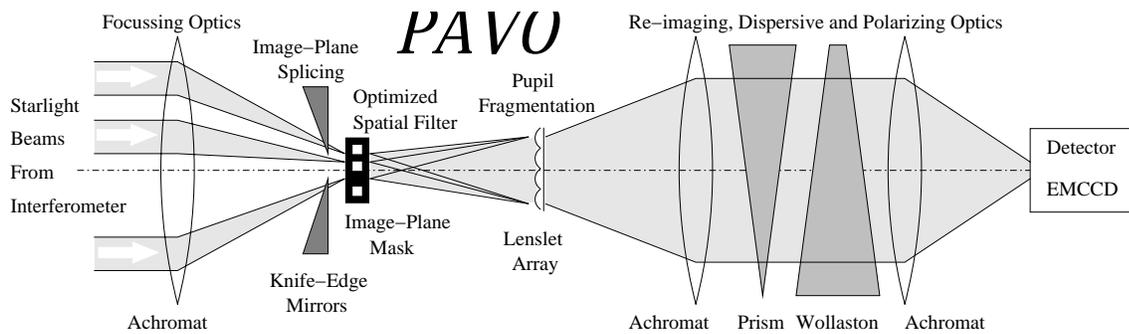


Figure 1. A schematic view of the PAVO optical concept. The Wollaston prism (or other polarizing optics) have yet to be implemented.

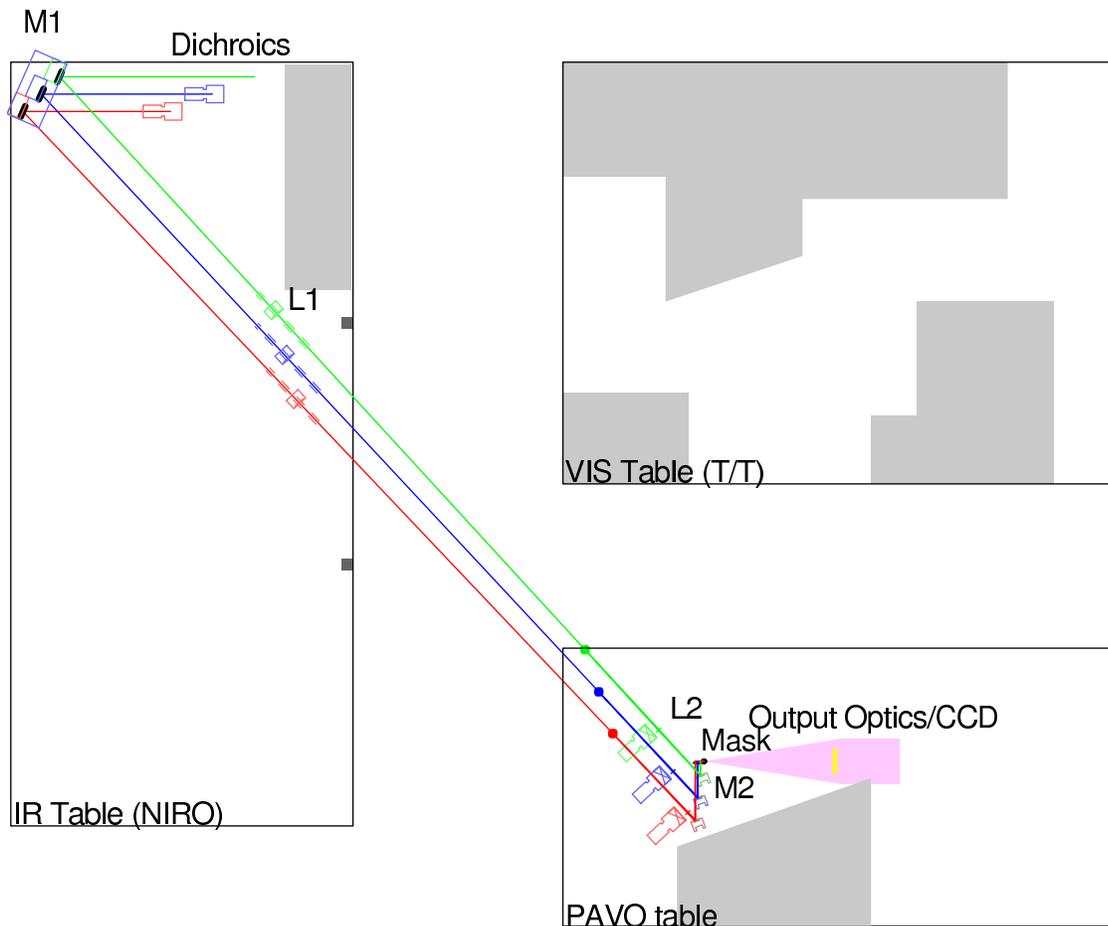


Figure 2. The overall layout of the PAVO beams in the CHARA optical lab. The left table is the CHARA Classic table, where the PAVO beams are extracted by near-normal incidence dichroics that reflect the beams upwards, ~ 100 mm above the standard CHARA beam height. This table also contains the 2 m focal length lenses. The main part of PAVO is located on one corner of the CHARA artificial light source table, in the lower-right of the diagram. The upper right optical table contains the tip/tilt system.

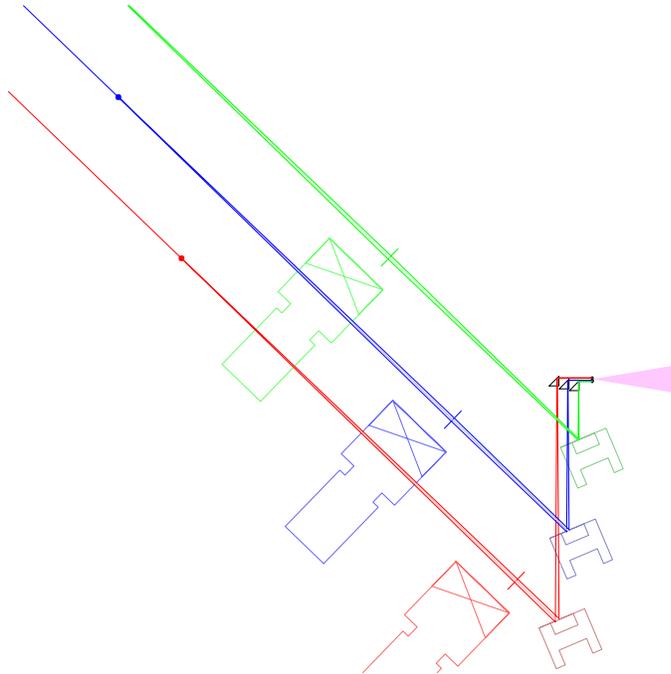


Figure 3. The layout of the lenses L2, the mirrors M2, the right-angle prisms and the mask of PAVO at CHARA. The three prisms are positioned along a line in a left-right direction but are placed on machined steps, so that they have different locations in a direction into the page. In this way, the beam from the left-most prism passes over the other two prisms, and the beam from the center prism passes over the right-most prism. The mask is placed approximately 2 mm to the right of the right-most prism.

The optical layout of PAVO at CHARA is shown in Figure 2, and a zoomed up portion of this layout around the mask is shown in Figure 3. This design brings star images from 3 separate $f/67$ beams side-by-side so that they each form an 8 mm wide image of the 1 m CHARA telescope pupil on a cylindrical lenslet array. As the CHARA pupils are at different distances depending on the telescope being used and the position of the delay lines, a nominal (average) pupil distance is used. The three beams first reflect off a dichroic then a 2 inch flat mirror before passing through a 2 lens imaging system. These widely-separated beams are brought together in two stages. Firstly, fold mirrors M2 that take the ~ 100 mm separation beams and place them 10 mm apart, using mounts that have one open side. Next, right-angle prisms utilizing total internal reflection steer the beams to be as little as 0.56 mm apart as they pass through the three-hole mask.

The optics beyond the 3-hole mask (shown schematically in Figure 1) consist of a filter wheel, one 0.5 mm pitch, 2.5 mm focal-length microlens array from SUSS micro-optics, two six-element 50 mm lenses from Schneider, and an anti-reflection coated 45-degree BK7 prism from Melles Griot. The lenslet array can slide out, being replaced with a 200 mm focal length lens and a narrow-band filter to directly image the mask. The camera is a back-illuminated DU-897 EMCCD from Andor, which uses 512x512 sensor from E2V. This camera is capable of photon-counting, and reads out at 100 Hz in a 1x4 binning mode as used by PAVO.

Figure 4 illustrates various aspects of the PAVO design with simulated fringes. Fringes from two telescopes at a single wavelength will form fringes like those displayed in the left figure. Seeing has caused the fringes to be curved and have non-uniform intensity across the pupil, but visibility is still essentially unity. The only causes of reduction in system fringe visibility here are differing intensity patterns across the two input pupils, effects of non-zero pixel size in sampling these fringes, temporal fringe blurring and unmatched Jones matrices in the two arms of the interferometer.

The right hand side of Figure 4 shows the simulated 2-telescope fringes as seen by the detector, after the pupil has been fragmented by the lenslet array and then dispersed by the prism. The splitting of light into two orthogonal polarization states as illustrated in Figure 1 has not yet been implemented. The near-vertical fringes

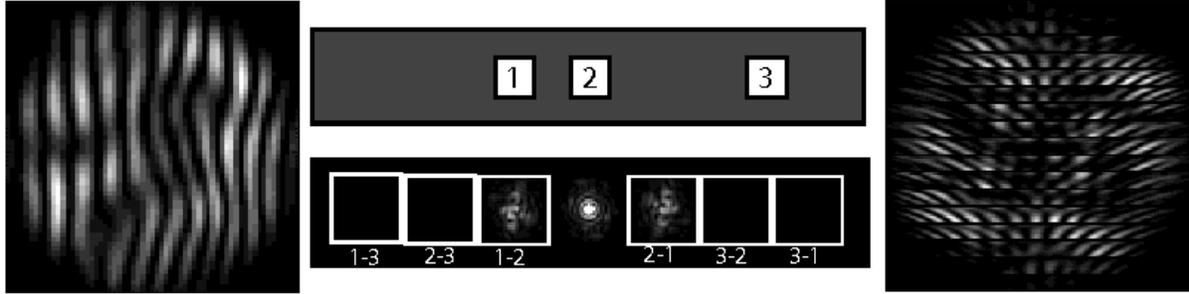


Figure 4. Left: Simulated fringes formed by a single wavelength at the front of the lenslet array. Right: Simulated fringes after passing through the lenslet array, neglecting image-plane distortion by the prism. The dispersion can be seen as an increase in fringe spacing as we move from the bottom to the top of each lenslet image. Center top: the non-redundant image-plane mask. Center bottom: The power spectrum of the simulated 2-telescope fringes, showing the distinct regions where fringes from each baseline appear. Note that vertical in a lab frame is to the right in all these diagrams (i.e. not up/down on the page).

from each lenslet at the center of the pupil illustrates that this simulated frame occurs at near zero group delay. Non-zero group delay results in the fringes within each dispersed lenslet image becoming tilted.

The center-top panel of Figure 4 shows the three hole mask. The holes are square so that they can have the maximum possible area and still be fully non-redundant. Fringe power (i.e. V^2 times squared flux) on each baseline can be found by taking the Fourier transform of the pupil-plane fringes. The squared-modulus of this Fourier transform is shown as a grey scale in the lower center panel. The fringes from each baseline are bounded by the non-overlapping square boxes in this panel, which correspond to non-zero regions of the mask transmission auto-correlation.

3. OPTIMIZING SENSITIVITY

One of the most important parts of a sensitivity optimized beam-combiner is the choice of wavelength. The general sensitivity limit for PAVO is taken to be the limit for visible group-delay tracking rather than the limit for infrared fringe tracking with PAVO collecting data passively. Putting practical considerations (e.g. instrumental noise, the availability of an infrared fringe-tracker) aside, we will first examine why this is a reasonable idea in principle. Sensitivity limits for optical interferometry have been examined several times before (e.g. Ref 4), but here we will ask a specific question applicable to PAVO, “What is the wavelength-dependence for the group-delay tracking limit on point-sources in the Rayleigh-Jeans limit (i.e. hot stars) with a fixed multi- r_0 aperture size?”.

Integration of group delay occurs both coherently and incoherently in two dimensions of space, one of time and one of wavelength. The temporal dimension is the simplest: our target for group-delay tracking will be a given signal-to-noise in group delay S_g per coherence time t_0 . We will ignore the effects of incoherent integration for group delay tracking, as this relates to the probability of tracking loss in a complex way that must include the effects of the outer scale and spectral dispersion is complicated by spectral resolution.

The number of photons received per coherence time per unit wavelength goes as:

$$\frac{dN_p}{d\lambda} \propto \lambda^{-9/5}. \quad (1)$$

This exponent includes a factor of -4 from the Rayleigh-Jeans tail, a factor of 1 from the number of photons per unit energy, and a factor of 6/5 from the coherence time. We will consider the limit at low signal-to-noise (S/N) per coherence volume, where signal-to-noise is proportional to N_p . The signal-to-noise per coherence patch goes as the square of the coherence length, and the number of patches goes as the inverse of the coherence length squared. These M patches increases the S/N by \sqrt{M} over the S/N for a single patch. So we have, for the S/N at fixed photon rate:

$$SN \propto \lambda^{6/5}. \quad (2)$$

This means that signal-to-noise at fixed bandwidth goes as:

$$SN_{\Delta\lambda} \propto \lambda^{-3/5}, \quad (3)$$

and signal-to-noise at fixed fractional bandwidth goes as:

$$SN_{\Delta\lambda/\lambda} \propto \lambda^{2/5}. \quad (4)$$

This expression is, however, only valid for small fractional bandwidths, where signal-to-noise is proportional to bandwidth. Once a bandwidth reaches a size where r_0 at a wavelength equal to the coherence length is smaller than the aperture-diameter, different wavelength regions can only add incoherently. The aperture diameter at which this occurs goes as:

$$D \propto \left(\frac{\lambda^2}{\Delta\lambda}\right)^{6/5}. \quad (5)$$

In turn, this means that the limiting bandwidth for coherent addition of different wavelength regions goes as:

$$\Delta\lambda_C \propto \lambda^{-2}. \quad (6)$$

Therefore, Equation 3 only applies to a bandwidth of $\Delta\lambda_C$, after which point signal-to-noise goes as the square root of $\Delta\lambda/\Delta\lambda_C$. In this regime, the exponent in Equation 3 should change to 2/5. For CHARA, when comparing R/I, H and K bands for a 1 m aperture diameter, all band-passes have a similar ($\sim 0.35 \mu\text{m}$) bandwidth, and we are in-between the regime where the exponent of Equation 3 is -3/5 and 2/5. Therefore, R/I, H and K bands should in principle have similar group-delay tracking performance.

In practice, there are many other factors that influence group-delay signal-to-noise in interferometers, such as thermal noise and detector readout noise. It is difficult to achieve photon-limited performance in the H- or K-bands*, and photon-counting at high quantum efficiency is only available for silicon-based detectors. Therefore, we consider PAVO to have a requirement that it is not only a science combiner to be used with infrared tracking (e.g. for sufficiently red or well-resolved stars), but should also provide sensitive group-delay tracking.

4. SPATIAL FILTERING

Previous attempts at spatial filtering have aimed at measuring a fringe signal in a single-mode, using single-mode fibers. Light from a multi- r_0 aperture is intrinsically multi-mode, so to take advantage of all the light, one wants to ideally measure the different modes separately, to gain a \sqrt{N} advantage in signal-to-noise for measuring N modes. One way to do this is to spatially-modulate fringes by placing the incoming beams side-by-side, then at least Nyquist sample the fringes on a detector in a conjugate plane. The VEGA/REGAIN beam combiner would accomplish this with an infinitely-narrow slit, but in practice a finite slit width in that instrument decreases fringe visibility because of wavefront aberrations in a direction orthogonal to the fringe direction.

The PAVO beam combiner concept reverses the image and pupil planes from a VEGA/REGAIN-like concept: placing the star images beside each other and forming spatially-modulated fringes in a pupil plane. The criteria for Nyquist sampling for fringes at a single wavelength λ is:

*For a 1 m aperture-size and 5% throughput to a 1 arcsec squared field-of-view, K-band is in principle photon-limited up to a magnitude of 8.5, but in practice photon-limited performance is not achieved

$$\Delta x_P < \frac{\lambda}{2\theta_m} \quad (7)$$

$$\approx \frac{\lambda(\mu\text{m})}{\theta_m(\text{''})} \times 0.1 \text{ m.} \quad (8)$$

$$(9)$$

Here Δx_P is the spacing of sampling points in a pupil-plane, and θ_m is the maximum angular scale in the image plane. With an image-plane mask design as shown in Figure 4, with the on-sky angular width of a hole equal to θ_h , θ_m is equal to $7\theta_h$ in the fringe direction, and θ_h in the orthogonal direction.

For the case where the pupils from the two starlight beams have equal intensity, the system visibility for PAVO, as measured by infinitesimally small pixels and lenslets is in principle unity. However, the finite size of pixels and lenslets reduce the fringe visibility. This can be calculated in the Fourier transform of the fringe plane F :

$$A(x', y') = \int F(x, y) \exp(2\pi i \frac{xx' + yy'}{\lambda D}) dx dy \quad (10)$$

Here D is the distance between image-plane and lenslet and λ is the wavelength. x and y are linear dimensions running along an image of the fringes at a single wavelength (e.g. left hand side of Figure 4). This plane is labeled A because for observations of a point source it represents the auto-correlation of the image-plane electric field (i.e. at the mask) in the Fraunhofer limit (within a multiplicative constant). The power in the fringe plane (F) is conserved in the autocorrelation plane (A), and the signal in the A plane can be corrected for the non-zero pixel and lenslet size by the following simple formula:

$$A_{\text{cor}}(x', y') = A(x, y) / \text{sinc}(\frac{\pi x' S_x}{\lambda D}) / \text{sinc}(\frac{\pi y' S_y}{\lambda D}). \quad (11)$$

Here S_x is the size of the detector pixels (i.e. in distance units), S_y is the size of the lenslets.

In beam combination utilizing single-mode fibers, accurate visibility amplitudes are only possible through the use of photometric taps, or by rapidly calibrating the relative fluxes between each beam. For PAVO, the equivalent process to this is to measure the relative intensity in each beam.

The advantage of forming fringes in a pupil-plane is less sensitivity to seeing and misalignment, at least in principle. If local seeing were to change at one telescope, or one beam train to become misaligned, then the fringe visibility in an image-plane combination scheme would change. In a pupil-plane spatially-modulated scheme, only the fringe spacing and tilts change, preserving the visibility. In practice, this advantage is not fully realized because the limited pixel count available for modern EMCCDs drives the image-plane spatial filter to be relatively small ($\sim 1''$ on a side), which in turn means that the spatial filter translates phase aberrations on the input pupil into intensity variations on the pupil as seen by the lenslet array.

5. DATA ANALYSIS

The analysis of PAVO data is both done in real-time in a C programming language, and in `yorick` for post-processing. The following sections focus on the real-time data analysis.

5.1 Group Delay Estimate

The data are recorded on a CCD that has 512 physical columns and rows. The parallel-clocking direction of the CCD is binned by a factor of 4, producing a logical 512×128 array, read out at 110 Hz. The raw frames show a quadratic image-plane distortion due to the prism. This effect is removed, and optionally a threshold is applied to the raw data when operating in a photon-counting mode.

Once a distortion solution has been applied, each column of the data array undergoes a 1-dimensional real to complex Fast-Fourier-Transform (FFT), using the FFTW library.⁵ Although the considerations of Section 4 were most intuitively realized with 2-dimensional FFTs, the symmetry of the array means that 1-Dimensional FFTs are adequate to split the data into fringes from each baseline. This FFT is represented by an operation on the undistorted image g_{jk} as:

$$s_{jk} = F_k(g_{jk}). \quad (12)$$

Here F_k represents a 1-dimensional Fourier transform in the k direction (down each column). We define the "mask size" m_j at column j to be equal to:

$$m_j = \frac{W_M W_P}{D_{M,L} \lambda_j} \quad (13)$$

Here W_P is the width of the pupil (in mm) at the lenslet array, W_M is the width of the mask hole (in mm) and $D_{M,L}$ is the distance between the mask and the lenslet array (also in mm). λ_j is the wavelength for column j . For baseline b (equal to 1, 2 or 3), the fringe center frequency is at $2bm_j$, and we define a window function $w_{b,jk}$:

$$w_{b,jk} = \begin{cases} 1 & (2b-1)m_j < k < (2b+1)m_j \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

Then the complex fringes in the pupil-plane $p_{b,jk}$ for baseline b is given by:

$$p_{b,jk} = F_k^{-1}(S_{-2bm_j}(F_k(g_{jk})w_{b,jk})). \quad (15)$$

Here S_{-2bm_j} represents a shift of the fringe signal by $-2bm_j$ units to the origin in the A -plane (c.f. Equation 10). For example, the baseline labeled 2-1 in the lower panel of Figure 4 is shifted one box to the left by this operation, and the power everywhere else is multiplied by zero by the $w_{b,jk}$ function. In the absence of aberrations, $p_{b,jk}$ would simply have to be integrated over the indices j and k to give a visibility estimate. However, due to aberrations and seeing, we have to treat $p_{b,jk}$ as having significant phase variation across the pupil.

The group delay estimate is constructed from phasors $q_{b,m}$ of the form:

$$q_{b,m} = \sum_j \sum_k p_{b,jk} p_{b,j(k+m)}^* \quad (16)$$

Here $*$ represents complex conjugation. m runs from 1 to 6 (see below), although it could in principle run higher. The sum is only evaluated over the high signal-to-noise part of each lenslet image (i.e. spectrum) and never in-between (where phase is discontinuous). For $m = 1$, the phase of $q_{b,m}$ is the phase difference between successive wavelength channels. This phase difference is equal to 2π at the formal edge of the coherence envelope of a single wavelength channel (i.e. assuming perfect optics), which is $\lambda^2/\Delta\lambda$, or, $1/\Delta\kappa$, with κ being the wave-number. As the wavelength channels for PAVO are roughly evenly spaced in wave-number, this form of the group delay estimate is equal in a noiseless case for all wavelength channels j : hence the integral over j . Simply taking $l \text{Arg}(q_{b,1})/2\pi$, with l the coherence length of a single channel, would be one estimator of group delay.

However, to maximize signal-to-noise for group delay, it is important to use values of m that are higher than 1. For example, $q_{b,2}$ makes two revolutions in the complex plane while $q_{b,1}$ only makes one, halving the error on

group delay if both $q_{b,m}$ values have the same signal-to-noise. Using values of m that are too high adds noise, as the approximation $l_j = \text{Constant}$ is no longer good enough, and effects of dispersion and speckle decorrelation mean that the parts of the speckle pattern at widely separated wavelengths that make it through the image plane mask can no longer be meaningfully compared. A maximum value of $m = 6$, covering about a quarter of our full bandwidth, was chosen to maximize signal-to-noise in typical observing conditions.

Values of $q_{b,m}$ are combined to give a group delay $G(\delta)$ estimate by a discrete Fourier transform:

$$G(\delta) = \sum q_{b,m} e^{2\pi i m \delta}. \quad (17)$$

Here δ represents normalized group-delay. The real part of $G(\delta)$ is maximized, and the group-delay estimate is given by $l\delta_{\text{max}}$.

A problem with this technique comes at group delays with moduli between $l/2$ and l . The discrete Fourier transform only runs between $-l/2$ and $l/2$, meaning that the estimator wraps-around and other information is required to determine the sign of the group delay. Although no additional information is so far used for group delay tracking, the algorithm performs adequately on-sky thus far.

5.2 V^2 Estimate

The main estimator for V^2 takes the quantities $q_{b,0}$ as defined in Equation 16, divides by squared flux and then normalizes. V^2 estimators for sub-bandpasses simply restrict the integration in the j (row)-direction for both flux and $q_{b,0}$. This estimator requires background subtraction in the numerator and denominator, as well as a photon-bias subtraction in the numerator.

A coherent estimator for V^2 is also formed from successively recorded frames:

$$V_{C,b}^2 \propto \frac{\sum_t \sum_j \sum_k \Re(p_{b,jk,t}(t) p_{b,jk,t+1}^*)}{\sum_t \sum_j \sum_k g_{jk,t}}. \quad (18)$$

Here we have explicitly included the integration over a time index t , and the symbol \Re denotes taking the real part of a complex number. The coherent estimator has the advantage that there is no bias in the estimator, and the ratio of the coherent and incoherent V^2 estimators is used to directly estimate t_0 and correct visibility reduction due to coherence time effects.

6. COMMISSIONING DATA

PAVO at CHARA was commissioned in April 2008, successfully tracking fringes on many objects, including HD 95884 at an R magnitude of 7.1, and has tracked fringes on 3 baselines simultaneously. So far, only single-baseline data has been analyzed, so closure-phase performance remains uncharacterized.

Figure 5 shows the system V^2 over a sub-bandpass of PAVO during a 50 minute time period where two visits were made to each of 4 stars. The effects of partially resolving the four stars has been taken into account. The RMS V^2 is somewhere between 4 and 8%, depending on whether some of the V^2 scatter between stars is astrophysical or merely systematic.

The rather low system visibility of $\sim 0.6-0.7$ could be caused by several factors. The correction of Equation 11 has not been applied to these data, which would increase system visibility by $\sim 10\%$. It could be further increased by $\sim 5\%$ due to combined effects of the broad bandwidth in each wavelength channel and group delay tracking errors. Any remaining system visibility losses must come from a combination of vibrations above 110 Hz and a mismatch in Jones matrices between the two interferometer arms (which seems unlikely to be a large effect given the near-symmetry of the CHARA beam trains).

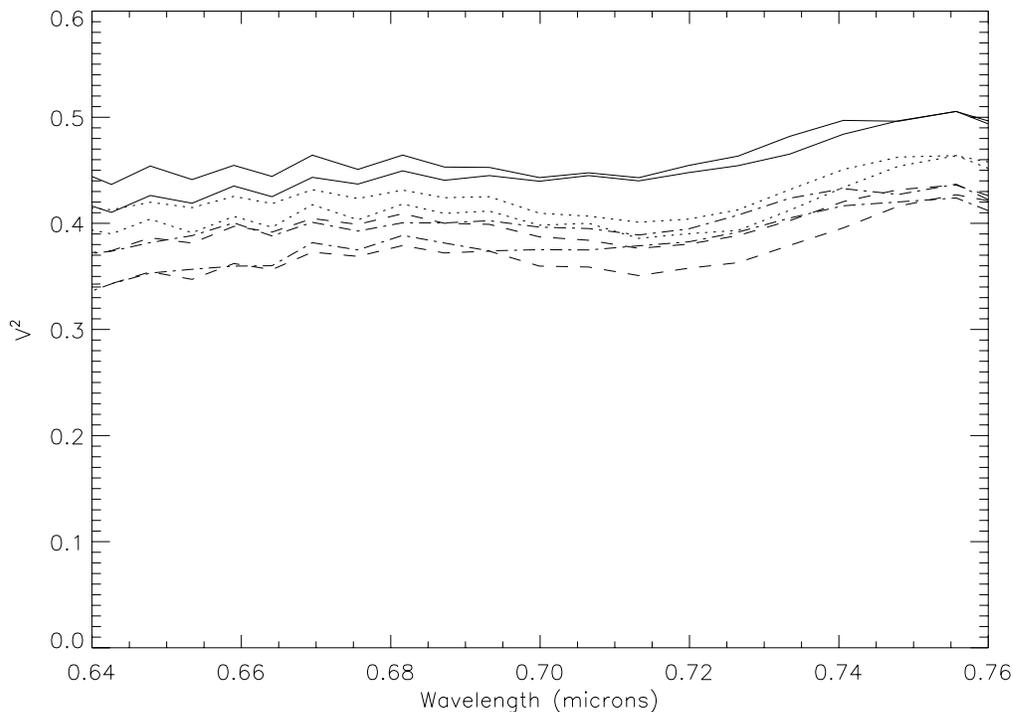


Figure 5. System V^2 derived from two separate observations of 4 stars: HD 106661 (solid lines), HD 107966 (dotted lines), HD 108382 (dashed lines) and HD 108765 (dot-dashed lines). These stars were observed on the W1W2 baseline of CHARA and have magnitudes between 4.9 and 5.7. The long-wavelength portion of the bandpass is not plotted, partially due to systematics arising from achromatism in the re-imaging optics (to be fixed for the July run).

7. CONCLUSIONS AND FUTURE WORK

The PAVO beam combiner is a spectrally-dispersed beam-combiner at CHARA designed to maximize sensitivity. Fringes are spatially-modulated and formed in a pupil plane, with a cylindrical-lenslet based integral field unit to sample the fringe pattern over a $\sim 40\%$ bandpass over the R and I bands. An image-plane mask band limits the fringe-power in the pupil plane so that all allowed modes of the fringe pattern are sensed, providing multi-mode spatial filtering. Commissioning data taken in April 2008 already demonstrates calibration at the $<5\%$ level in V^2 , and the ability to track fringes from three baselines simultaneously. The first science run will be in July 2008.

ACKNOWLEDGMENTS

M.I. would like to acknowledge Michelson Fellowship support from the Michelson Science Center and the NASA Navigator Program. A.M. acknowledges support from the W.M. Keck foundation. PAVO at CHARA has also been funded by the W.M. Keck foundation. The CHARA night assistants P.J. Goldfinger and Chris Farrington assisted greatly in making the commissioning observations possible.

REFERENCES

- [1] Davis, J., Ireland, M. J., Chow, J., Jacob, A. P., Lucas, R. E., North, J. R., O'Byrne, J. W., Owens, S. M., Robertson, J. G., Seneta, E. B., Tango, W. J., and Tuthill, P. G., "The Sydney University Stellar Interferometer: A Major Upgrade to Spectral Coverage and Performance," *Publications of the Astronomical Society of Australia* **24**, 138–150 (Oct. 2007).

- [2] Armstrong, J. T., Mozurkewich, D., Rickard, L. J., Hutter, D. J., Benson, J. A., Bowers, P. F., Elias, N. M., Hummel, C. A., Johnston, K. J., Buscher, D. F., Clark, J. H., Ha, L., Ling, L.-C., White, N. M., and Simon, R. S., "The Navy Prototype Optical Interferometer," *ApJ* **496**, 550 (Mar. 1998).
- [3] Mourard, D., Bonneau, D., Clausse, J.-M., Hénault, F., Marcotto, A., Blazit, A., Bosio, S., Bresson, Y., ten Brummelaar, T., Kervella, P., Lagarde, S., McAlister, H. A., Mérand, A., Merlin, G., Nardetto, N., Petrov, R., Roussel, A., Rousselet-Perraut, K., Stee, P., Sturmman, J., Sturmman, L., and Tallon-Bosc, I., "VEGA: a visible spectrograph and polarimeter for CHARA," in [*Advances in Stellar Interferometry. Edited by Monnier, John D.; Schöller, Markus; Danchi, William C. Proceedings of the SPIE, Volume 6268, pp. 62683Q (2006).*], Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference **6268** (July 2006).
- [4] Buscher, D., "Optimizing a ground-based optical interferometer for sensitivity at low light levels," *MNRAS* **235**, 1203 (1988).
- [5] Frigo, M. and Johnson, S. G., "The design and implementation of FFTW3," *Proceedings of the IEEE* **93**(2), 216–231 (2005). special issue on "Program Generation, Optimization, and Platform Adaptation".