LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Iteratively Weighted Centroiding for Shack-Hartmann Wave-Front Sensors

K. L. Baker, M. M. Moallem

April 5, 2007

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Iteratively weighted centroiding for Shack-Hartmann wave-front sensors

K.L. Baker

Lawrence Livermore National Laboratory, Livermore, CA, USA
Baker7@llnl.gov

M.M. Moallem

KM Optix, Fremont, CA, USA
MmoallemOD@yahoo.com

Abstract: Several techniques have been used with Shack-Hartmann wave-front sensors to determine the local wave-front gradient across each lenslet. In this article we introduce an iterative weighted technique which is specifically targeted for open-loop applications such as aberrometers and metrology. In this article the iterative centroiding technique is compared to existing techniques such as center-of-mass with thresholding, weighted center-of-gravity, matched filter and cross-correlation. Under conditions of low signal-to-noise ratio, the iterative weighted centroiding algorithm is demonstrated to produce a lower variance in the reconstructed phase than existing techniques. The iteratively weighted algorithm was also compared in closed-loop and demonstrated to have the lowest error variance along with the weighted center-of-gravity, however, the iteratively weighted algorithm removes the bulk of the aberration in roughly half the iterations than the weighted center-of-gravity algorithm. This iterative weighted algorithm is also well suited to applications such as guiding on telescopes.

OCIS codes: (010.1080) Adaptive optics; (010.7350) Wave-front sensing.

References and links

1. John W. Hardy, *Adaptive Optics for Astronomical Telescopes*. (Oxford University Press, Oxford, 1998).
2. Jeffrey A. Koch, Robert W. Presta, Richard A. Sacks, Richard A. Zacharias, Erlan S. Bliss, Michael J. Dailey, Mark Feldman, Andrew A. Grey, Fred R. Holdener, Joseph T. Salmon, Lynn G. Seppala, John S. Toeppen, Lewis Van Atta, Bruno M. Van Wonterghem, Wayne T. Whistler, Scott E. Winters, and Bruce W. Woods, "Experimental comparison of a Shack-Hartmann sensor and a phase-shifting interferometer for large-optics metrology applications," Appl. Opt. **39,** 4540 (2000).
3. Bernd Schafer and Klaus Mann, "Investigation of the propagation characteristics of excimer lasers using a Hartmann-Shack sensor," Rev. Sci. Instrum. **71,** 2663 (2000).
4. K. L. Baker, J. Brase, M. Kartz, S. S. Olivier, B. Sawvel, and J. Tucker, "The use of a Shack–Hartmann wave front sensor for electron density characterization of high density plasmas," Rev. Sci. Instrum. **73**, 3784 (2002).
5. N. Qi, R. R. Prasad, K. Campbell, P. Coleman, M. Krishnan, B. V. Weber, S. J. Stephanakis, and D. Mosher, "Laser wave-front analyzer for imploding plasma density and current profile measurements," Rev. Sci. Instrum. **75**, 3442 (2004).
6. J. Arines and J. Ares, "Minimum variance centroid thresholding," Opt. Lett. **27**, 497 (2002).
7. M. Nicolle, T. Fusco, G. Rousset, and V. Michau, "Improvement of Shack-Hartmann wave-front sensor measurement for extreme adaptive optics," Opt. Lett. **29**, 2743 (2004).
8. T. Fusco, M. Nicolle, G. Rousset, V. Michau, J.-L. Beuzit, and D. Mouillet, "Optimisation of a Shack-Hartmann-based wave-front sensor for XAO system," SPIE **5490**, 1155 (2004).
9. T. R. Rimmele, "Solar Adaptive Optics," SPIE **4007**, 218 (2000).
10. Lisa A. Poyneer, David W. Palmer, Kai N. LaFortune, and Brian Bauman, "Experimental results for correlation-based wave-front sensing," SPIE **5894**, 58940N (2005).
11. Troy Rhoadarmer, "Laboratory Demonstration of a Correlation-Based AO System for Wave Front Sensing of Extended Objects," AMOS Technical Conference, Maui, HI, 2006.
12. V.N. Dvornychenko, "Bounds on Deterministic Correlation Functions with Applications to Registration," IEEE Trans. Pattern Anal. Mach. Intell. **5**, 206 (1983).

13. Qi Tian and Michael N. Huhns, "Algorithms for Subpixel Registration," Comput. Vision Graph. Image Process. **35**, 220 (1986).
14. Yuval Carmon and Erez N. Ribak, "Phase retrieval by demodulation of a Hartmann–Shack sensor," Opt. Commun. **215**, 285 (2003).
15. Sabrina Velghe, Jérôme Primot, Nicolas Guérineau, Mathieu Cohen, and Benoit Wattellier,, "Wave-front reconstruction from multidirectional phase derivatives generated by multilateral shearing interferometers," Opt. Lett. **30**, 245 (2005).
16. Malgorzata Kujawinska, in *Interferogram Analysis*, edited by David W Robinson and Graeme T Reid (Institute of Physics Publishing, Bristol and Philadelphia, 1993).
17. Yuval Carmon and Erez N. Ribak, "Fast Fourier demodulation," Appl. Phys. Lett. **84**, 4656 (2004).
18. Pedro M. Prieto, Fernando Vargas-Martin, Stefan Goelz and Pablo Artal, "Analysis of the performance of the Hartmann-Shack sensor in the human eye," J. Opt. Soc. Am. A **17**, 1388 (2000).
19. Nick Kaiser, Gordon Squires, and Tom Broadhurst, "A Method For Weak Lensing Observations," ApJ **449**, 460 (1995).
20. Wim de Vries, Scot Olivier, Steve Asztalos, Leslie Rosenberg, and Kevin Baker, "Image Ellipticity From Atmospheric Aberrations," Publ. Astron. Soc. Of the Pacific (to be submitted in 2007).
21. Scot S. Olivier, Steven M. Jones, Diana C. Chen, Robert J. Zawadzki, Stacey S. Choi, Sophie P. Laut, and John S. Werner, "Biomedical Optics: OCT sees the human retina sharply with adaptive optics," Laser Focus World **42**, 89 (2006).
22. Amos Talmi and Erez N. Ribak, "Direct demodulation of Hartmann-Shack patterns," J. Opt. Soc. Am. A **21**, 632 (2004).
23. J.-L. Starck, F. Murtagh, and A. Bijaoui, *Image Processing and Data Analysis: The Multiscale Approach.* (Cambridge University Press, 1998).
24. K.L. Baker, "Least-Squares Wave-Front Reconstruction of Shack-Hartmann Sensors and Shearing Interferometers using Multigrid Techniques," Rev. Sci. Instrum. **76**, 053502 (2005).

1. Introduction

Shack-Hartmann wave-front sensors are the most prolific wave-fronts sensors currently in use in closed-loop adaptive optics systems.[1] Their popularity is also growing in a number of metrology applications and in the case of aberrometers, used to measure aberrations in the eye, they represent the foundation for the majority of the commercially available instruments. Other metrology applications for which Shack-Hartmann wave-front sensors are also utilized include; optical metrology[2], diagnosing aberrations present in lasers[3], wafer inspection and in determining density and magnetic field profiles in plasmas.[4,5] In these applications, the Shack-Hartmann wave-front sensor is not being used in a closed-loop adaptive optics system but rather as a stand alone measurement of the wave-front. The aberrations measured by these systems can be large and noisy. Often times in aberrometers a large number of CCD pixels are used for every subaperture in the lenslet array, which can lead to extremely noisy measurements due to read noise, dark current and scattered light. Each of these noise terms wreaks havoc with conventional centre-of-mass centroiding techniques. In this article current algorithms used to determine the wave-front gradients across the Shack-Hartmann lenslet array are reviewed. Techniques such as noise filtering and, in the case of matched filtering and correlation sensing, utilizing simulated references to improve the performance of these algorithms is examined. An iterative weighted algorithm is then presented which, as is shown in later sections, is particularly well suited for open-loop measurements of aberrated beams. A comparison of the iteratively weighted algorithm with existing algorithms is then presented both in open-loop, as well as in closed-loop. The results are then summarized in the last section of the article.

2. Existing algorithms

The most basic and widely used technique for determining the gradients from the spot pattern produced by a Shack-Hartmann lenslet array is the centre-of-mass (COM) centroiding algorithm. This technique is expressed mathematically in equations 2 a-c and 3 below with the weighting function, w(x,y), set to unity. The x and y centroids are simply the sum of the intensity pattern multiplied by x and y, respectively, divided by the sum of the intensity pattern. The problems that arise with this technique are due to the read noise, dark current and background light associated with each of the pixels on the wave-front sensor camera. As is readily apparent from Eq. 2 below, the pixels far away from the centre are multiplied by a larger number. As a consequence, the noise contribution far from the centre can quickly corrupt or overwhelm the actual signal when a large number of detector pixels, greater than the minimum four pixels, are used.

One technique that improves upon the centre-of-mass algorithm when noise is present is to utilize thresholding (T-COM) prior to the application of the COM algorithm.[6] This thresholding process identifies the noise level, subtracts the noise level from the subaperture array and sets all of the resulting negative pixels to zero. This algorithm can be expressed by Eq. 2 a-c and Eq. 3 below with the intensity, I(x,y), in these equations replaced by the intensity minus the noise level. The weighting function in these equations is replaced by a binary mask with the location of all positive pixels(I – noise level) set to 1 and all negative pixels set to 0.

Another variation that has recently been proposed is to implement the centroiding algorithm with Gaussian weighting to overcome the limitations due to noise.[7,8] This technique has been called weighted centre-of-gravity (WCOG). In this technique, the authors use a Gaussian weighting function, which is fully characterized by its full-width-at-half-maximum (FWHM), to determine the local gradients in each of the subapertures. This would then correspond to equations 1, 2 a-c and 3 below, however, the Gaussian weighting function of Eq. 1 would not have the ability change its centroid location, x_c=0 and y_c=0. By choosing the FWHM of the Gaussian weighting function to be less than the size of the subaperture, the noise contribution from those pixels far from the reference spot is reduced significantly. The application of Gaussian weighting was demonstrated by these authors to greatly reduce the effects of noise on determining the local gradients.

A separate technique to determine the wave-front gradients was developed by the solar physics community which came to be known as correlation sensing.[9] The solar physics community utilizes Shack-Hartmann wave-front sensors in closed-loop adaptive optics systems to look at an extended object, the sun. For this application researchers needed a technique that would allow the gradients in each of the subapertures to be determined but they could not utilize the conventional centre-of-mass technique. Their solution was to use cross-correlation to determine the local gradients in each of the subapertures. Cross-correlation is performed between a given subaperture and all of the remaining subapertures to determine the local displacement of the spots in each of the subapertues, i.e. the displacement of the reference image that provides the maximum correlation. This technique has also been applied to the post-analysis of open-loop intracavity laser and telescope adaptive optics systems.[10] Analogous to this technique is matched filtering, which is very similar to correlation sensing. In this case a convolution in the spatial domain is performed and the peak of the resultant signal is used to determine the gradients in a given subaperture.[11] For both of these techniques the subpixel shift is determined by fitting the correlation peak using parabolic interpolation.[12,13] In this process a three-by-three subarray is extracted from the correlation array with the correlation peak located at the central pixel. The x and y subpixel shifts are given by the x and y derivative(second order) of the correlation peak divided by its x and y Laplacian at the central pixel, respectively. For a symmetric object, these techniques should perform equivalently.

A separate technique that has recently been developed examines the similarity between Shack-Hartmann sensing and carrier frequency interferometry to determine the phase.[14,15,16] In

this case the spot pattern formed by the Hartmann sensor is treated like an interference pattern and Fourier transformed. The resulting spots to the right and left of the central "dc" spot have encoded the x gradients and the spots above and below the central "dc" spot contain the y gradient information. By shifting one of the x spots to the centre and filtering out all the information outside of a circle placed about this spot, an inverse Fourier transform returns the wrapped x gradient. Repeating this process with one of the y spots returns the wrapped y gradient. By performing an unwrapping procedure, the gradients at each of the subapertures are determined. In a regime where the aberrations are small, this approach has been further simplified.[17]

An iterative windowing technique has also been developed to reduce the noise in situations where a large number of pixels is contained within a given subaperture.[18] This windowing technique estimates the centroid by conventional centre-of-mass techniques and then decreases the window size by one pixel, catered about the estimated centroid location. The authors continue this process until the window size is approximately the theoretical size of the first lobe of the diffraction pattern. This technique is equivalent to Eq. 2 a-c and Eq. 3 below with the weighting function equal to one within the area of the window and zero outside of the window. The window is then iteratively decreased in size and kept centred on the previous centre-of-mass estimate.

3. Iterative weighted algorithm

The basis of this algorithm finds its roots in the weak lensing astrophysics community. In this case the desire is to determine the ellipticity correlations between galaxies in an effort to map out the dark matter in the universe. One such weak lensing data pipeline uses a Gaussian weighting function which is characterized by its FWHM and by its centroid location, both of which can be iteratively adjusted.[19,20] From this weighted distribution, the moments, up to second order, of the intensities in each of the subapertures can be found. The x and y centroids are simply the sum of the intensity pattern multiplied by the weighting function and x and y, respectively, divided by the sum of the intensity pattern multiplied by the weighting function. These can then be iteratively applied to the centre-of-mass of the original weighting Gaussian function to refine the location of the centroid of the Gaussian and minimize the noise. In addition, using the higher order moments of the intensity pattern provides an estimation of the FWHM of the Gaussian which can also be iterated upon along with the centroid of the Gaussian. This then forms the basis for an iterative scheme which sets the centroid and FWHM of the weighting function in the presence of noise. In the limit that the ability to iteratively change the location of the centroid and its FWHM is removed, this method reduces to the weighted centre-of-gravity (WCOG) algorithm described above. By allowing the iterative process to occur, however, the algorithm is able to more accurately locate the centroids in cases of large aberrations where the Hartmann spots are not close to the reference spots and noise is present. In addition, this technique could be used to more quickly bring the AO system to the null operating point when operating in closed-loop as discussed below.

The weighting function, $w(x, y)$, is expressed by a Gaussian distribution centred at (x_c, y_c) with a full-width-at-half-maximum equal to $2\sigma\mathrm{SQRT}(2\ln2)$. The weighting function is given by

$$w(x, y) = \frac{1}{2\pi\pi^2} e^{-\frac{1}{2\sigma^2}(x-x_c)^2} e^{-\frac{1}{2\sigma^2}(y-y_c)^2} , \qquad (1)$$

with σ representing the root-mean-square (RMS) deviation of position from the mean. This definition of the weighting function results in a normalization of the integral over all space of unity. For the simulations, this integral is approximated by a sum of the weighting function

over the given subaperture. In this expression x, x_c, y, y_c and σ represent real numbers. The various position moments, up to second order, are expressed as

$$S_w = \sum_{x,y} I(x,\ y)w(x,\ y),$$ (2a)

$$S_x = \sum_{x,y} xI(x,\ y)w(x,\ y),$$ (2b)

$$S_y = \sum_{x,y} yI(x,\ y)w(x,\ y),$$ (2c)

$$S_{xx} = \sum_{x,y} x^2 I(x,\ y)w(x,\ y),$$ (2d)

$$S_{yy} = \sum_{x,y} y^2 I(x,\ y)w(x,\ y),$$ (2e)

where S_w, S_x, S_y, S_{xx} and S_{yy} represent the different moments of the lenslet spot pattern, I(x,y), on the CCD camera with the weighting function. The average location of the Gaussian, or its centroid, is simply

$$x_c = S_x/S_w, \quad y_c = S_y/S_w$$ (3)

where x_c and y_c represent the x and y centroid locations. For a Gaussian spot, the RMS deviation of position from the mean corresponds to the weighted $(x-x_c)^2 + (y-y_c)^2 = \sigma^2$ or

$$\sigma = \left(\frac{1}{S_w}\right)\sqrt{S_{xx}S_w - S_xS_x + S_{yy}S_w - S_yS_y}$$ (4)

This set of equations form the basis for the iteratively weighted centroiding algorithm in which the centroid location and the FWHM of the weighting function can be adaptively changed.

As pointed out by the referee, this technique could also be utilized using a recentred average of the Hartmann spots as the weighting function rather than a Gaussian or Sinc[2] function. This technique is also applicable to extended scenes as in the case of correlation sensing. For this choice of the weighting function the FWHM would not be changed such that Eqs. 2 a-c and 3 would be iterated upon with the weighting function consisting of a normalized average of the recentred Hartmann images/spots. For the simulations presented in this article, however, a Gaussian weighting function was used and Eqs. 1-4 were utilized for the iterative weighted algorithm. It is also the case that σ can be fixed and not iterated upon with a Gaussian or Sinc[2] weighting function as well.

To investigate the convergence of the iteratively weighted algorithm, numerical wave optics simulations were performed. The simulations consisted of a lenslet array with 16 lenslets across the diameter of a circular pupil. The phase aberration used to test the iteratively weighted algorithm was based on a scaled version of Zernike coefficients measured in a human eye.[21] The Hartmann spots were then formed by Fourier transforming the electric field at each of the lenslets in the array and multiplying by the complex conjugate to obtain the intensities. Both photon noise and read noise were then introduced to the focal plane containing the Hartmann spots. The iteratively weighted algorithm was used to estimate the slopes which were then reconstructed. The variance between the original and reconstructed

phases was then calculated. The variance in the reconstructed wavefront, as a function of the number of iterations was then used to test the convergence of the iteratively weighted algorithm. Figure 1 shows the residual variance as a function of iterations through Eqs. 1-4. This figure shows that the algorithm converges within about four steps for this initial aberration(peak-to-valley ~ 15 radians). Based on this result, all of the simulations in this article which use the iteratively weighted algorithm were performed with four iterations.
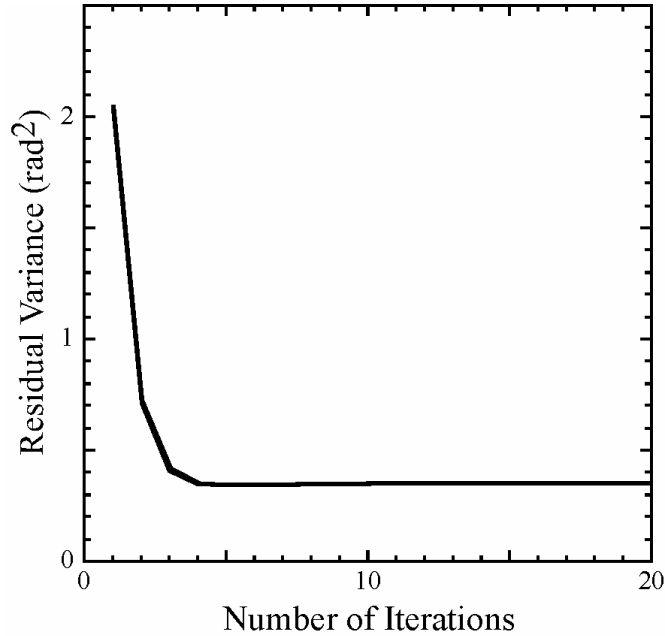

Fig. 1. Residual variance as a function of the number of iterations performed with the iteratively weighted algorithm.

The numerical simulations can also be utilized to test the width of the Gaussian weighting function after the four iterations as a function of the initial guess. Figure 2 shows the width of the Gaussian weighting function, σ, after four iterations, solid red line, as a function of the initial starting value, σ_{init}, for the Gaussian width. The spread in starting values for σ ranged over a factor of close to 5, from 1.26 to 6 pixels. The resultant values for the weighting functions, σ, after four iterations varied by less than seven percent. The simulated Hartmann spots had a full-width-at-half-maximum of 3.8 pixels which corresponds to a Gaussian σ_s of 1.6 pixels. Figure 2 also shows the reconstructed wavefront variance, solid black line, using four iterations of the iterative weighted algorithm as a function of the initial value for the weighting function, σ_{init}. The weighting function σ that resulted in the lowest reconstructed wavefront variance was $\sigma = 1.62$ pixels. This illustrates for this particular phase profile that over a broad range of initial starting values of the weighting function, σ_{init}, both larger and smaller, that after four iterations the algorithm converges to a σ value that is within several percent of the Hartmann spot full-width-at-half-maximum. With this particular phase profile, an initial starting value, $\sigma_{init} = 2.25$ pixels, that was slightly larger than the spot size, $\sigma_s = 1.6$ pixels, resulted in the lowest variance.

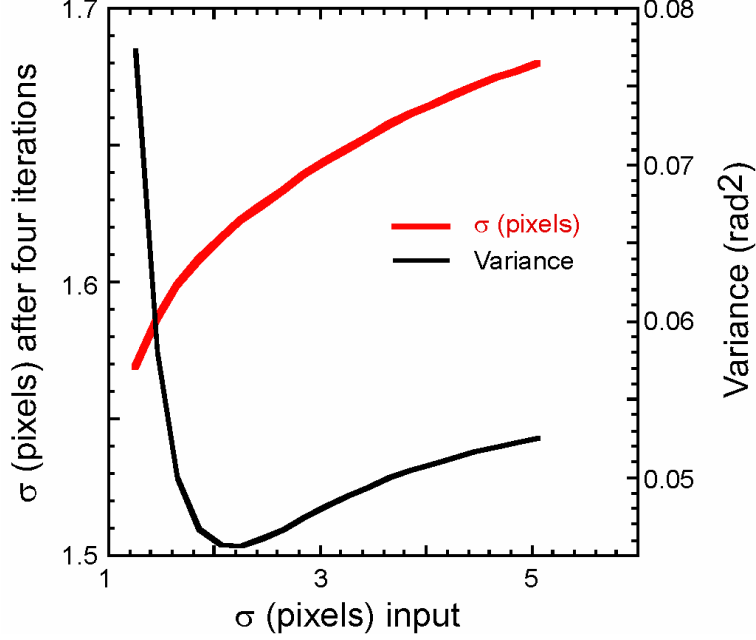Fig. 2. Weighting function σ and reconstructed wavefront variance after four iterations as a function of the initial starting value, σ_{init}, for the Gaussian width.

The iteratively weighted algorithm is very effective at placing the weighting function in the correct location. Once the weighting function is correctly positioned over the Hartmann spot, then the noise characteristics of the iteratively weighted algorithm are identical to the weighted centre-of-gravity algorithm discussed above. The authors of the WCOG technique[7,8] analyzed the noise properties of using a Gaussian weighting function and derived the following two expressions for the noise variance in the case of photon noise, σ_{ph}^2, and detector noise, σ_{det}^2, as

$$\sigma_{ph}^2 = \frac{\pi^2}{2ln(2)\mathcal{N}_{ph}}\left(\frac{N_T}{N_D}\right)^2\left(\frac{N_T^2 + N_W^2}{2N_T^2 + N_W^2}\right)^2 \text{ and} \tag{5a}$$

$$\sigma_{det}^2 = \frac{\pi^3}{32[ln(2)]^2}\left(\frac{\sigma_{detector}}{\mathcal{N}_{ph}}\right)^2\left(\frac{N_T^2 + N_W^2}{N_D}\right)^2, \tag{5b}$$

where $\mathcal{N}_{ph}$ is the number of photons, N_W is the FWHM of the Gaussian weighting function, N_T is the Full Width at Half Maximum (FWHM) of the seeing-limited Point Spread Function (PSF) in pixels, N_D is the FWHM of the diffraction limited point spread function and $\sigma_{detector}$ is the standard deviation of the detector noise.

The iteratively weighted algorithm could also be applied to the functions of field acquisition and guiding on astronomical telescopes. The weighting function in Eq. 1 could also be different in each of the subapertures to include effects such as laser guide star parallax for the next generation of large diameter astronomical telescopes. A more general expression for the Gaussian weighting function would then be expressed as $w(x,y)=(2\pi\sigma_x\sigma_y)^{-1}EXP\{-[(x-x_c)cos\theta-(y-y_c)sin\theta]^2/(2\sigma_x^2)-[(x-x_c)sin\theta+(y-y_c)cos\theta]^2/(2\sigma_y^2)\}$ with σ_x and σ_y representing the root-mean-square (RMS) deviation of position from the mean along the x and y axes, respectively. The angle θ represents the rotation of the Gaussian from the x-axis. For the results presented in this article, a symmetric weighting function was assumed such that $\theta = 0$

degrees and $\sigma_x = \sigma_y = \sigma$. In this case, the more general weighting function reduces to the form given in Eq. 1 above.

4.   Noise filtering and reference generation

Nearly all of the centroiding techniques, except for the conventional centre-of-mass algorithm, perform better after filtering when the spots formed by the lenslet array are larger than a single pixel. In this case, the noise, which varies at the spatial scale of an individual pixel, can be reduced by filtering the signal prior to analysis by the above mentioned gradient determining algorithms without affecting the signal itself. This is not the case when the spot size from the lenslet array is chosen to be a single pixel in width such that its spatial variation is comparable with the background noise. In this case filtering will affect both the noise and the signal you are trying to detect. For this article, where the spot size was larger than a single pixel, several filtering techniques were investigated for their utility in reducing the noise and improving the reconstructed wave-front variance. The techniques examined included filtering with wavelet transforms, Fourier transforms and spatial convolution with Gaussians. In general the wavelet transform resulted in the lowest residual variance followed by Fourier transforms and then filtering using a spatial convolution with a Gaussian. Other authors have used filters previously including triangular and square convolutions in real space[22] and box car averaging in Fourier space,[14,22] which was required for this particular algorithm.

The wavelet transform that was evaluated for noise filtering was specifically the à trous algorithm, which is for use with discreet data.[23] The image to be filtered in this case is the intensity spot pattern recorded on the wave-front sensor due to the lenslet array. The different scales of the image are formed by taking convolutions of the image with a scaling function that maintains its discreet shape but doubles in spatial scale between successive wavelet scales or convolutions. A B_3-spine was chosen as the scaling function such that the coefficients of the convolution mask are given by the one-dimensional array, (1/16, 1/4, 3/8, 1/4, 1/16), for a vector and the two-dimensional array,

$$\begin{pmatrix} \dfrac{1}{16} & \dfrac{1}{4} & \dfrac{3}{8} & \dfrac{1}{4} & \dfrac{1}{16} \end{pmatrix} \otimes \begin{pmatrix} \dfrac{1}{16} \\ \dfrac{1}{4} \\ \dfrac{3}{8} \\ \dfrac{1}{4} \\ \dfrac{1}{16} \end{pmatrix}, \tag{6}$$

for images, where $\otimes$ represents the Kronecker product.[23] From the above description, it is readily apparent that the method of removing noise using spatial convolution with a Gaussian is quite similar to the implementation of the à trous wavelet transform. Fourier transforms were also investigated as a means of removing the high spatial frequencies associated with the detector noise. In this case the image was Fourier transformed and a circular binary mask applied to block out the high spatial frequencies before inverse Fourier transforming back. This is similar to the filtering method used in carrier frequency interferometry.[16] The spatial filtering in this article was applied to the focal plane containing all of the Hartmann spots, as opposed to each of the subapertures separately. This focal plane was larger than the pupil and

so no zero padding was utilized as the distance between the outer Hartmann spots and the edge of the focal plane was greater than the extent of the filters.

The three filtering schemes were tested on the reconstruction of a phase profile with a circular aperture. The read noise on the detector was set at 20 e- root-mean-square and the number of photons per subaperture was varied between 50 and 1000. The results, shown in Fig. 1 below, demonstrate the utility of filtering the signal before determining the gradients when the signal-to-noise-ratio decreases as the unfiltered signal, the solid green line, had the highest residual variance. The à trous wavelet transform, dashed black line, had slightly better performance, particularly at the lowest SNR, than the Fourier transform filter, solid blue line, which performed slightly better than the method utilizing convolution with Gaussians, dashed red line.
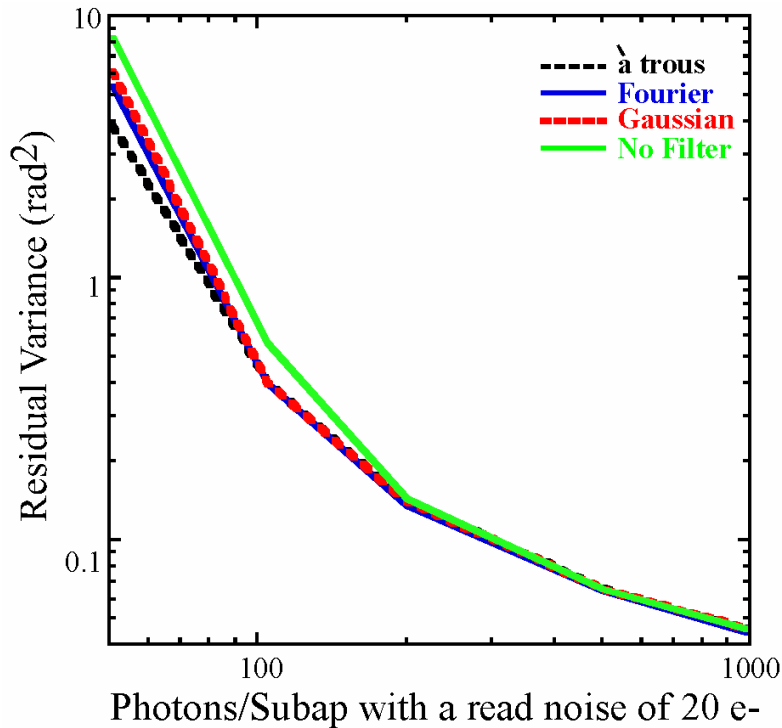

Fig. 3. Residual variance as a function of noise for the three different filtering techniques and the unfiltered case.

In addition to filtering the noise, centroid locating techniques such as correlation sensing and matched-filter also need a reference image. This reference can be taken from one of the subapertures, filtered or unfiltered, or it can be generated analytically without noise. This can lead to better performance for high noise scenarios and also enables the techniques to accurately reconstruct the tip/tilt component of the phase aberration, a feature which is compromised when comparing to one of the subapertures. If one of the Hartmann spots from the lenslet array is used as the reference, then that spot can still be compared to a simulated spot to recentre the reference before that spot is compared with all of the remaining spots in the lenslet array to allow determination of the tip/tilt information. Figure 4 below shows the residual variance as a function of photons/subap for the cases where the reference is simulated(dashed black line), taken from the noise filtered centre subaperture(solid blue line) and taken from the unfiltered centre subaperture(dashed red line). The intensity profile on the

CCD camera was filtered with the Fourier transform filter and the correlation sensing method was used to determine the centroids. The tip/tilt was also removed from the gradients before and after reconstruction to place all three reference generation methods on a comparable footing. As Fig. 4 shows, the simulated reference performed better than the technique of using one of the subapertures at lower SNRs. The overall variance is lower in Fig. 4, relative to Fig. 3, due to the subtraction of the tip/tilt aberration.
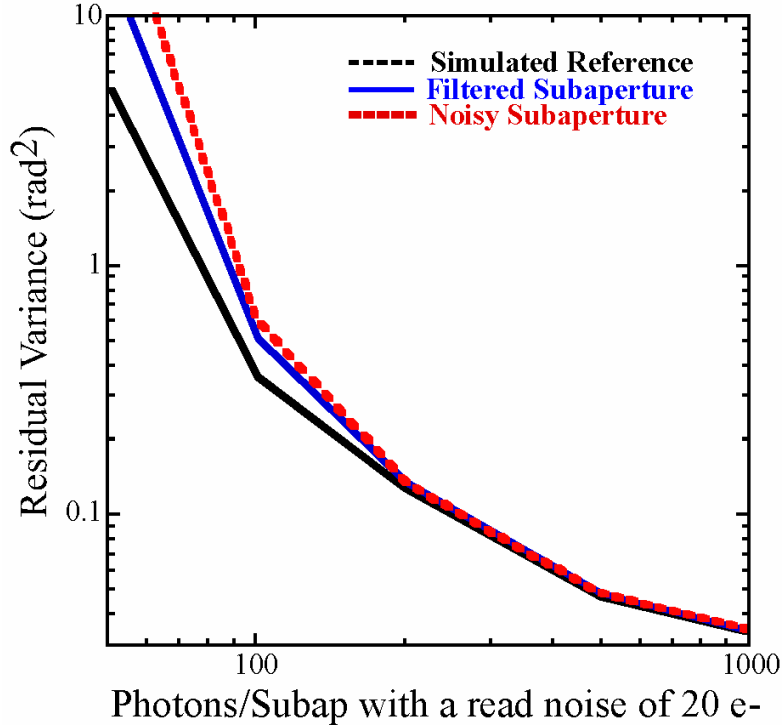

Fig. 4. Residual variance as a function of noise for the three different reference generation methods.

5. Comparison of algorithms

A comparison of the algorithms discussed above was performed. This comparison involved starting from an initial phase aberration, expressed as Zernike coefficients, at a circular aperture. At this plane a lenslet array was simulated such that there were a total of sixteen lenslets across the diameter of the circular aperture and the electric field across each of the lenslets was parsed into 32 by 32 phase elements. In addition, each lenslet represented 16 by 16 pixels on the CCD wave-front camera. The Hartmann spots were then formed by Fourier transforming the electric field at each of the lenslets in the array and multiplying by the complex conjugate to obtain the intensities. Both photon noise and read noise were then introduced to the focal plane containing the Hartmann spots. In all of the simulations the read noise was fixed at 20 e- rms and the number of photons-per-pixel was varied. The residual variance between the original phase and the reconstructed phase was then determined as a function of the number of photons-per-pixel. The different algorithms were then used to identify the displacement of the Hartmann spots caused by the phase aberration and a Multigrid reconstructor[24] was used to reconstruct the phase based on the local gradients determined by the different algorithms. The initial phase aberration used to test the various centroiding algorithms was based on a scaled version of Zernike coefficients measured in a human eye.[21] Through filtering the noise and utilizing a simulated reference spot, the cross-

correlation, matched filtering, centre-of-mass, windowed and centre-of-mass threshold algorithms detailed above were optimized before comparing them to the weighted centre-of-gravity and the weighted iterative algorithm expressed in Eqs.1-4.

The results from the various algorithms are shown below in Fig. 5. At the higher photons-per-subaperture end of the graph, all of the algorithms except the weighted centre-of-gravity (WCOG) algorithm, solid black line, converge as expected. This algorithm is higher than the others simply because the aberration is large such that the spots are displaced from the nominal centre of the CCD pixels representing the lenlset array. To encompass the displaced spot, the WCOG algorithm must have a large FWHM Gaussian weighting function which makes it much more like the conventional centre-of-mass algorithm and thus very susceptible to noise. As shown below, when the WCOG algorithm is operated in closed-loop the spot displacement is brought to zero allowing a much smaller FWHM to be used. This enables the WCOG algorithm to achieve the best performance along with the iteratively weighted algorithm. As the number of photons-per-subaperture is decreased, the measurements become noisier and the remaining algorithms begin to diverge in their performance. After the WCOG algorithm, the threshold centre-of-mass (T-COG) algorithm, dashed blue line, and the windowed algorithm, solid red line, have the next largest variance. This is followed by the matched filter, solid green line, and cross-correlation algorithms, solid grey line. These two algorithms have virtually identical performance except at the noisiest point in the graph where the matched filter algorithm has a slightly higher variance than the cross-correlation algorithm. This is expected as the Hartmann spots are symmetric and both techniques use parabolic interpolation to achieve subpixel accuracy. The best performance is achieved by the iteratively weighted algorithm, dashed black line, presented in this article. The performance improvement over the existing algorithms is particularly high under conditions of low signal-to-noise ratio as can be seen in Fig. 5. As the number of photons decreases, the slope of the variance vs. photon number, N_{ph}, for the iterative weighted algorithm closely matches the $(1/N_{ph})^2$, dashed grey line, behaviour expected form Eq. 5b above.
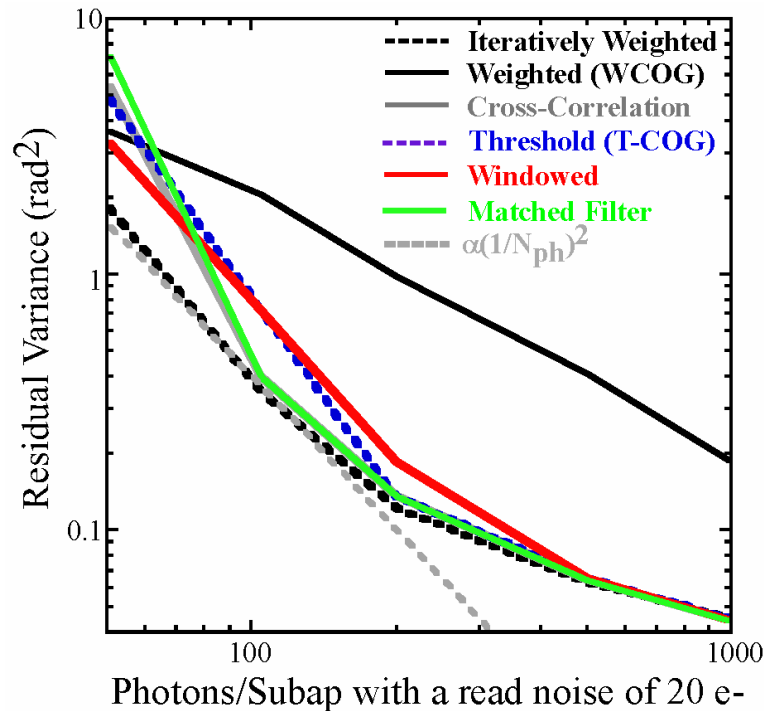

Fig. 5. Comparison of the different centroid locating algorithms in open-loop.

It is also of interest to determine how well these algorithms operate in closed-loop since nearly all adaptive optics(AO) systems operate in this manner and AO systems represent a large fraction of the Shack-Hartmann wave-front sensors in operation. The same phase aberration discussed above was used to test the closed-loop operation of the algorithms. The phase was reconstructed as detailed above, but a gain multiplier of 0.8 was applied to it before being subtracted from the original phase. A comparison of the iteratively weighted, WCOG, cross-correlation and centre-of-mass algorithms was undertaken and the results are shown below in Fig. 6. Figure 6a represents the lower signal-to-noise-ratio of the two figures at 105 photons-per-subaperture and 20 e- rms noise. The centre-of-mass, dashed grey line, exhibits the worst performance followed by an approximately equivalent performance by the cross-correlation algorithm, solid grey line, the threshold centre-of-mass algorithm, solid blue line, and the windowing algorithm, solid red line. The iteratively weighted, dashed black line, and WCOG algorithms, solid black line, converge to the same nominal level. The iterative functionality of the iterative weighted algorithm was turned off after the first three iterations to reduce the computational load. By incorporating the iterative capability in the first few times through the loop, however, the rate at which the aberrations are removed is accelerated considerably over the WCOG algorithm as seen in Fig. 6a and 6b. Figure 6b represents a higher signal-to-noise ratio such that there is less of a disparity between the cross-correlation algorithm and the iteratively weighted and WCOG algorithms. Again the iteratively weighted algorithm removes the bulk of the phase aberrations in roughly half the iterations through the loop, at gain=0.8, than the WCOG algorithm.
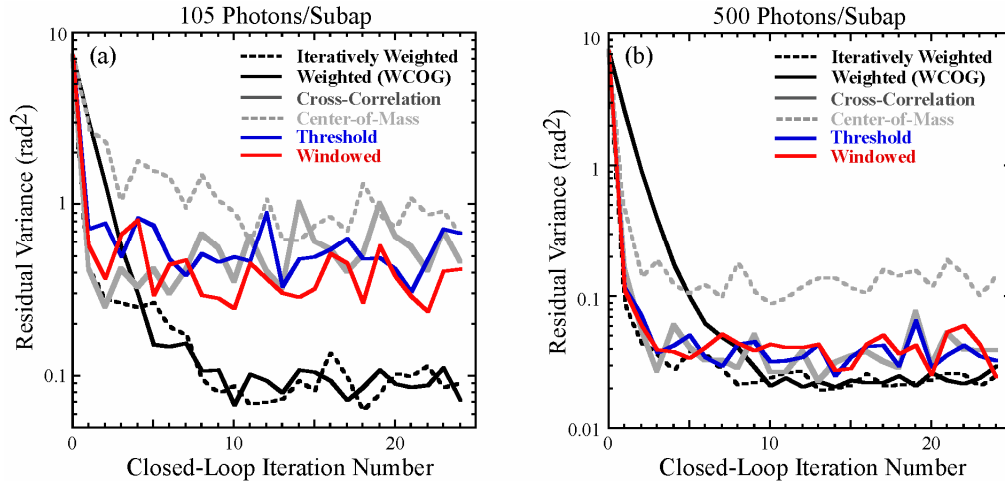

Fig. 6. Comparison between different algorithms in closed-loop. Figure 4a illustrates the results with 105 photons/sub-aperture and Fig. 4b represents the case of 500 photons/sub-aperture.

6. Summary

In this article an iteratively weighted centroiding algorithm was developed for wave-front sensors that measure the gradient of the wave-front, such as Shack-Hartmann and shearing interferometers. This algorithm was specifically designed to mitigate the deleterious effects of noise and is ideally suited for use in open-loop applications where the Hartmann spots are far from their unaberrated positions. A comparison between the iteratively weighted centroiding algorithm was performed with other commonly used centroiding algorithms including; center-of-mass, center-of-mass threshold, cross-correlation, matched filter, windowing and weighted

center-of-gravity. In order to provide the best comparison between the algorithms, methods such as filtering the noise on the simulated CCD wave-front camera and using simulated reference spots were investigated to improve the performance of algorithms such as center-of-mass threshold, cross-correlation and matched-filter. In the case of filtering the noise on the simulated CCD wave-front camera, Fourier filtering, filtering with à trous wavelet transform and filtering by Gaussian convolution were investigated. The best performance was achieved with the à trous wavelet transform, however, all three methods of filtering out-performed the unfiltered case. The filtering was performed for the case where the full-width-at-half-maximum of the Hartmann spots on the CCD wave-front camera was several pixels, 3.8, in width. The matched-filter and cross-correlation algorithms were also compared using simulated reference spots and using a spot from one of the lenslets as a reference, both noise filtered and unfiltered. Because the method of using one of the spots from a lenslet as a reference is more susceptible to tilt errors, the tilt was removed for these comparisons. The results indicated that using a simulated spot as the reference resulted in a lower error variance in the reconstructed wave-front than using a spot from one of the lenslets as a reference.

The performance of all of the algorithms compared was significantly better than the traditional center-of-mass algorithm in both open and closed-loop as expected. The open-loop comparison showed that the iteratively weighted algorithm achieved the lowest residual variance of all of the algorithms tested, as shown in Figs.5, under conditions of low signal-to-noise ratio. This was the case even when the optimizing techniques discussed above were used on the other algorithms. The iteratively weighted centroiding algorithm was also investigated in closed-loop operation, the most common mode of operation for adaptive optics systems. In this case the iteratively weighted and weighted center-of-gravity centroiding algorithms gave comparable levels of performance and resulted in a lower variance than the other algorithms. The iteratively weighted centroiding algorithm did, however, remove the bulk of the phase aberrations in roughly half the iterations than the WCOG algorithm.

Acknowledgements