Estimating Phase Errors from Pupil Discontinuities from Simulated On Sky Data: Examples with VLT and Keck

Masen Lamb[a,b], Carlos Correia[c], Jean-François Sauvage[c,d], Jean-Pierre Véran[b], David Andersen[b], Arthur Vigan[c], Peter Wizinowich[e], Marcos van Dam[f], Laurent Mugnier[d], Charlotte Bond[c]

[a]University of Victoria, 3800 Finnerty Rd, Victoria, Canada;
[b]NRC Herzberg Astronomy, 5071 W. Saanich Rd, Victoria, Canada;
[c]Aix Marseille Université, CNRS, LAM (Laboratoire d'Astrophysique de Marseille) UMR 7326, 13388, Marseille, France;
[d]ONERA, 29 Avenue de la Division Leclerc, 92320 Châtillon, France;
[e]W. M. Keck Observatory, 65-1120 Mamalahoa Hwy, HI 96743 Kamuela, USA;
[f]Flat Wavefronts, 21 Lascelles Street, Christchurch 8022, New Zealand

ABSTRACT

We propose and apply two methods for estimating phase discontinuities for two realistic scenarios on VLT and Keck. The methods use both phase diversity and a form of image sharpening. For the case of VLT, we simulate the 'low wind effect' (LWE) which is responsible for focal plane errors in low wind and good seeing conditions. We successfully estimate the LWE using both methods, and show that using both methods both independently and together yields promising results. We also show the use of single image phase diversity in the LWE estimation, and show that it too yields promising results. Finally, we simulate segmented piston effects on Keck/NIRC2 images and successfully recover the induced phase errors using single image phase diversity. We also show that on Keck we can estimate both the segmented piston errors and any Zernike modes affiliated with the non-common path.

Keywords: Phase Diversity, Image Sharpening, Wavefront Sensing, Adaptive Optics

1. INTRODUCTION

Piston discontinuities in segmented pupils are difficult to quantify when considering traditional AO systems; this mainly originates from the inability of the Shack-Hartmann wavefront sensor (WFS) to estimate differential piston. A prime example of this low wind effect (LWE) on the SPHERE/VLT system, where it is believed the relative lack of wind causes large temperature gradients across the VLT pupil. These temperature gradients are thought to exist in quadrants defined by the secondary struts, which act as thermal barriers between segments. The result is one quadrant will have a different index of refraction than another, specifically due the air's refractive index being dependant on temperature. A uniform temperature in one quadrant will have roughly a net piston effect when compared to a quadrant with a different uniform temperature. The Shack-Hartmann WFS is simply unable to detect differential piston and the resulting PSF reveals features akin to 'Mickey Mouse Ears'; correcting this effect is crucial considering the purpose of SPHERE is to achieve the highest contrast possible, which is clearly contaminated by this effect.

Similarly, segmented telescopes such as Keck can suffer from differential segment piston, which also has an effect on the observed PSF. These effects have been identified on Keck as 'low order residual errors';[1,2] the net result is a PSF with improper structure - mostly in the first Airy ring. Some of these errors may arise from the inability of the SHWFS to detect differential piston, however it has been shown that the SHWFS is indeed able to sense some of these low order residuals.[3]

Solutions to these two pupil effects have been investigated using a variety of different approaches. For example recent studies (N'Diaye et al. 2016, submitted) have suggested the use of a Zernike WFS (ZELDA[4]), capable

Further author information: M.L.: E-mail: masen@uvic.ca

of detecting these piston variations, which could be used to estimate the LWE on SPHERE. In the case of the Keck low order residuals, work has been done by van Dam 2016[3] employing, for example a Gerchberg-Saxton approach to the phase estimation. Furthermore, a technique known as Phase Diversity has also been explored (a description of this technique will be addressed later in this paper) considering long exposure images through turbulence, both on sky and in simulation.[2,3] In this paper we explore two different methods for correcting these effects for on-sky situations. The first method is the well-known approach of Phase Diversity,[5] where known diverse images are compared to synthetic images to estimate the phase of an optical system. The second method is the technique of Focal Plane Sharpening,[6] where the PSF in the focal plane is optimized using a deformable mirror. We simulate realistic images on both the VLT/SPHERE system and the Keck/NIRC2 system and investigate the feasibility of these methods in correcting both. We also explore the concept of single image Phase Diversity, which could be very useful when trying to avoid evolving seeing and AO residuals between a set of on-sky images. Finally, the technique of Focal Plane Sharpening (FPS) has no model limitations if the PSF is optimized over the peak of the PSF (for example), whereas Phase Diversity relies on a model for the creation of synthetic images, and any errors associated with that model will effect its performance. We therefore hypothesize that FPS will out-perform Phase Diversity given its model-free quality, and investigate this quantitatively in the case of the LWE estimation on SPHERE.

2. ESTIMATION METHODS

2.1 Phase Diversity

Our Phase Diversity code uses the approach of Paxman et al. (1992),[5] where an aberration-only objective function and its gradient are fed through a non-linear optimization algorithm to minimize the quadratic difference between synthetic and data (from the OTFs of the images). The aberration-only objective function consists of the coefficients of the basis to be estimated (i.e. Zernikes, or any other type of basis). The stopping criterion is defined by a tolerance parameter (which we choose as 1e-9, see our companion SPIE paper (no. 9909-247) for more information) - when the quadratic difference of the OTFs is below this value we reach convergence. The synthetic and real data images typically contain in and out of focus images, although the code can incorporate any number of images or type of diversity. The optimization the algorithm we employ is the quasi-Newton BroydenFletcherGoldfarbShanno (BFGS) algorithm. The code has the ability to either jointly estimate the object along with the phase, or to just estimate the phase itself and assume the object known (which we simplify as a point source). The Phase Diversity code has been developed as class for the AO MATLAB software OOMAO.[7] For more details of the code see the companion SPIE paper (no. 9909-247), and for an excellent overview of the Phase Diversity technique we refer the reader to Mugnier et al. (2006).[8] We adopt a 10% error for each mode used in the creation of our diverse images to simulate realistic errors.

2.2 Focal Plane Sharpening

Focal Plane Sharpening has also been developed as a class for OOMAO,[7] and has cross-compatibility with its Phase Diversity counter part. The algorithm receives the focal plane PSF as input and optimizes on a variety of criteria (chosen by the user) using the Nelder-Mead downhill-simplex method;[9] the input parameters to the optimization method are the basis used to create the PSF (which can be DM influence functions, Zernike Modes, or any combination of modes chosen by the user). The method along with criteria choices are summarized in detail in the companion SPIE paper (no. 9909-247). For this paper, at each step in the optimization we: use a small field of view centred on the PSF, median-filter the image with a 2*FWHM box to reduce noise, and fit a 2D-gaussian and measure the amplitude. This amplitude is optimized for its peak value, changing at each iteration with the new set of basis coefficients; in this sense FPS may be considered an 'on-line' optimization whereas Phase Diversity would be 'off-line'.

3. CASE STUDY 1: ESTIMATING THE LOW WIND EFFECT ON SPHERE

3.1 Basis and simulated images

To estimate the low wind effect on SPHERE we propose a basis roughly defining the pupil plane phase variations that occur at each quadrant of the VLT pupil (see Figure 1, top). The basic principle of our analysis is as follows:

estimate the coefficients of this modal basis using both phase diversity and focal plane sharpening and assess their qualities. The amplitude of the piston errors we use to artificially produce the LWE are about 1200 nm peak-to-valley (PV) and are shown in Figure 1 (bottom-right); this LWE error is conservative as typical LWE errors exist in the range of 600-800 nm PV WFE. We also adopt typical non-common path aberrations (NCPA) representive of the SPHERE system, specifically 40 nm rms WFE following a $1/\nu^2$ power law; they are also shown in the same Figure.

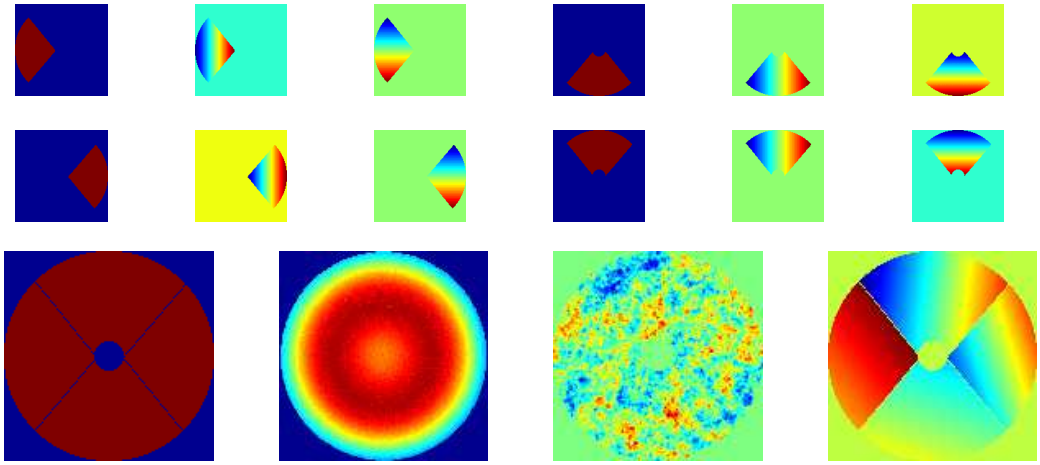

Figure 1. Top: Piston, tip, and tilt basis used to recreate the PSF variations seen during the low wind effect on SPHERE. Each mode is normalized to 1 rad rms (except the pistons). Bottom from left to right: VLT pupil, SPHERE apodization mask, assumed NCPA corresponding to 40 nm rms WFE, and 1200 nm PV WFE low wind effect errors.

To estimate the coefficients of our basis both our phase diversity and FPS algorithms require images of the PSF. We wish to represent the DTTS imager on SPHERE as accurately as possible in the creation of these images. We therefore simulate H-band images that we assume are dominated by photon and read noise. The images are created using the adaptive optics simulation tool OOMAO. The images are created with a 32x32 pupil sampling. Adding photon and read-noise is straightforward with this software, and we adopt a read-noise of 10e, as is typical for a Hawaii I detector. We simulate a DM with 41 actuators across the pupil. Images are created with a sampling as close to the DTTS imager as possible ($\sim$ 4 pixels across the FWHM). We apply an AO correction with this DM to a turbulence phase map generated from an r_0 of 11 cm (at 550 nm) at a sampling time of 500 Hz. Long exposure images are created by assuming a windspeed (15 m/s) and stepping the turbulence over the corresponding number of sampling times during the integration of the images. This is particularly important since the DTTS imager typically acquires $\sim$1 second exposures. We note however that no residual AO phase errors are incorporated in the generation of these long exposure images (i.e. lag, aliasing, etc.) and we will have this long exposure functionality in the near future. However we are currently able to apply these residual phase errors to instantaneous images, and we generate these in a different analysis of Keck images in this paper. Finally, the images used with both phase diversity and focal plane sharpening contain a field of view within the correctable region of the DM. This is not extremely important for focal plane sharpening, but for phase diversity it is extremely important: we find we have serious errors otherwise, increasing in effect as more of the uncorrected halo contaminates the image. Furthermore, the diversity we choose (i.e. focus) is always aimed to have its intensity distribution contained within the 'dark' correctible region.

3.2 Estimation methods

3.2.1 Phase Diversity

For the estimation of the LWE modal coefficients we employ what we consider 'traditional' Phase Diversity - when two images are used with focus as the diversity and the object is simultaneously estimated. Due to the combination of 40 nm rms NCPA and the relatively large PV amplitude of the LWE ($\sim$1200 PV nm), we consider two waves PV of focus for our diversity.

We consider the scenario where a star's signal to noise is comparable with actual data, to do this we consider the on sky data shown in Figure 2. We adopt a SNR of a typical star taken from this data set (SNR $\sim$ 70). Using a star with this SNR, we find we have difficulty on convergence using phase diversity, where half the time the solution will 'run away' in tip and tilt and converge on an erroneous result. We consider 3 approaches to solving this issue:

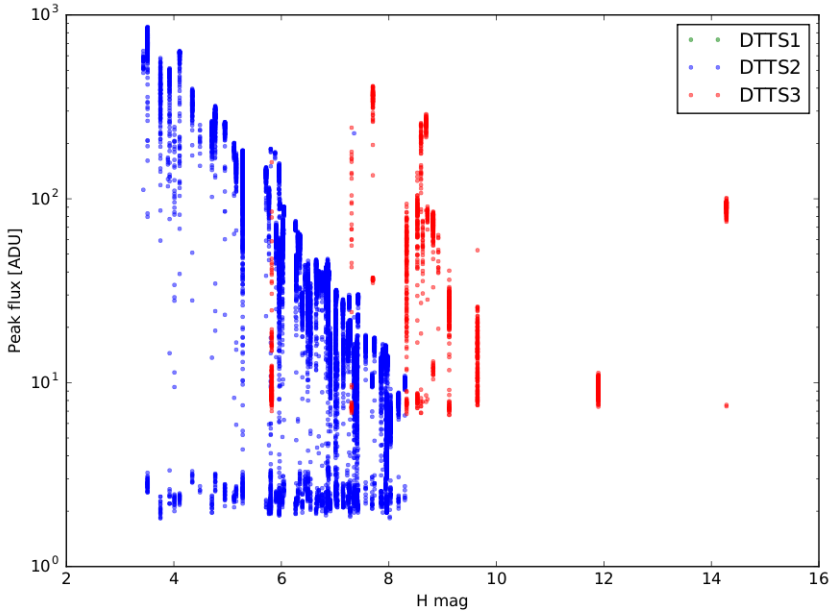

Figure 2. DTTS imager data, taken from SPHERE. The different colors correspond to different acquisition modes. The cloud of points around 2-3 ADU correspond to mis-detections, and we take this as the noise. Note: the values in this curve are subject to the inherent 20 nm PV focus on the DTTS imager, which results in a lower peak intensity than the true data shown here. After considering the noise floor and the data points adjusted for the 20 nm focus, we estimate a typical star has a SNR $\sim$ 70, and use this value for our analysis.

1. Increase the SNR of the image, presumably by exposing for longer, however on the real system this will integrate residual AO effects such as lag error, etc. We find this solution works well (see Figure 3), however we do not simulate these long exposure errors and caution that these errors should be considered for a more in-depth analysis. Given the practical simplicity of this solution, it seems like a viable realistic option.

2. The convergence seems to break down because of a signal-to-noise issue, therefore it is worth considering a different type of diverse image to see if it performs better under noise. Higher diversity modes, specifically cousins of the trefoil family (i.e. Z_{11}, Z_{21}, etc.), seem to work better in our simulation. In Particular we consider Z_{66}, where we have chosen this mode because it is a relatively high mode - which we find produce better phase estimates - and it is not too high such that it will be difficult to create with a SPHERE-like DM. When this mode is used in our simulation it always produces a better estimate of the phase (by a few nm rms), and it converges faster than the focus-diversity case. We adopt 2 waves (PV) of Z_{66} as our diversity and our results are shown in Figure 3 (bottom left image). One problem with realistically implementing this approach could arise from the high order phase speckles introduced from this diversity, which may be hard to disentangle from realistic AO phase residuals and high order, uncorrected NCPA. Furthermore, creating higher mode shapes such as Z_{66} will always have a higher residual error than creating a simpler mode such as focus.

3. Assume the object is known, which should be reasonable considering a star is effectively a point source, thus simplifying the phase diversity algorithm. The immediate result is that the estimate is not as accurate,

however it seems to have a much faster convergence rate. Figure 3 shows results from our simulation using this method.

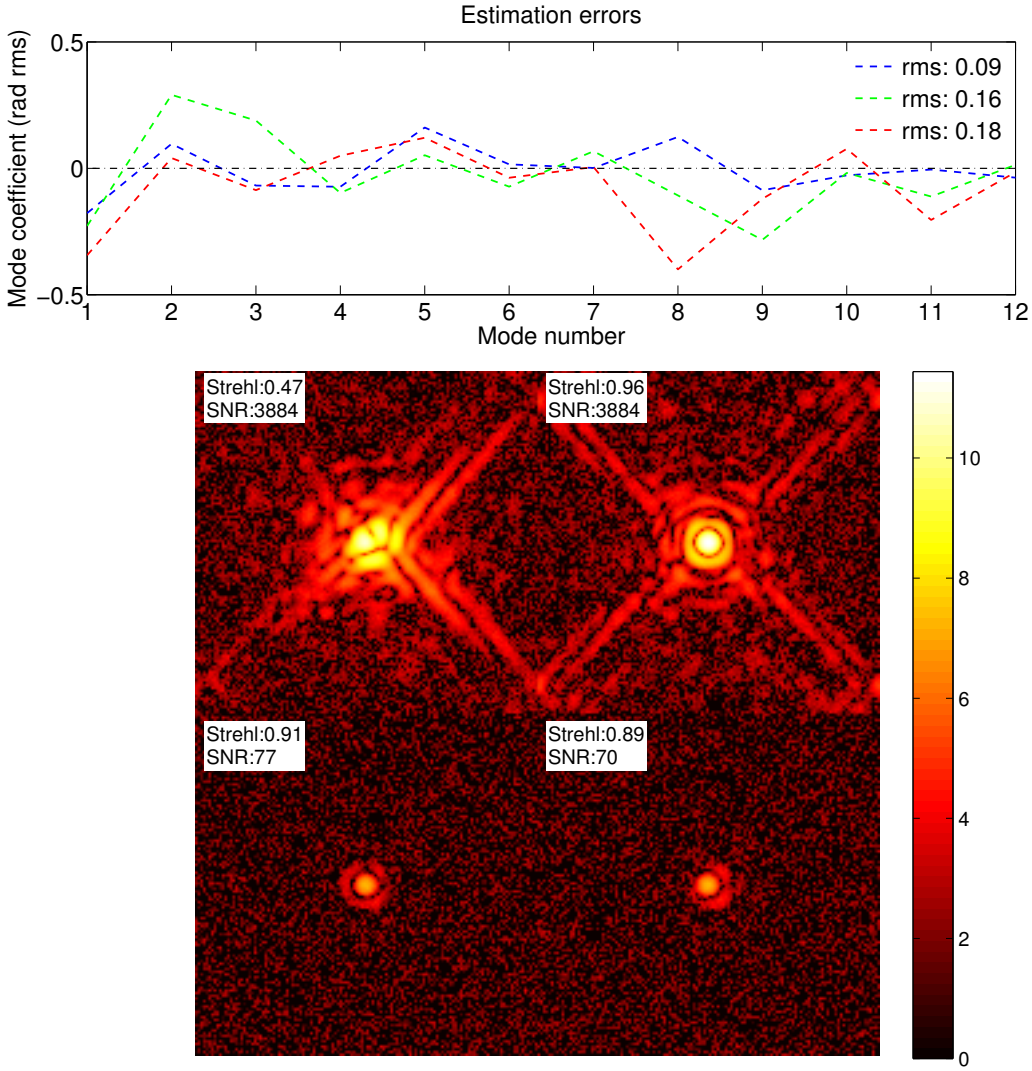

Figure 3. Top: residuals of LWE piston, tip, and tilt estimations from the actual modes, using phase diversity for 3 different scenarios (blue: long exposure object-estimation with focus diversity, green: object-estimation with higher diversity (Z_{66}), and blue: assumed point source with focus-diversity). It can be seen the long exposure (blue) scenario performs the best, as indicated its rms residual from the actual modes. Bottom: simulated VLT images, created from the phase projection of the estimated modes. Top left/right: no correction/long exposure. Bottom left/right: higher diversity/assumed point source. From these images, the highest performance in terms of Strehl is clearly using the long exposure image. The bottom two images have diffraction rings that fall under the pedestal of the noise.

From all three of these scenarios we conclude

- Under typical DTTS imaging conditions, traditional phase diversity does not reliably work. If longer exposures do not contain too many adverse AO phase residuals, then phase diversity with long exposure images solves this issue.
- Otherwise: using typical DTTS SNR values, phase diversity can work with a higher order diversity, or with assuming a known object. In the former case the estimate is more accurate, and in the latter the speed is considerably faster.

Table 1. Phase Diversity and Focal Plane Sharpening results correcting for the Low Wind Effect.

Case	Method	Resid. coeff. rms (rad)	Resid. wavefront** rms (nm)	Strehl (Marechal)	No. images	Initial SNR
Case 1	PD-1 (assumed object)	0.17	89.4	89.1	2	70
	PD-2 (long exposure)*	0.09	55.3	95.7	2	3.8e4
Case 2	FPS	0.40	56.1	95.1	122	106
Case 3	FPS + PD-1	0.12	57.9	95.2	64+2	119
Case 4	FPS + PD-2	0.11	55.7	95.6	30+2	129
Case 5	FPS low SNR	0.28	75.4	91.3	121	10

*Shown for reference
**The actual wavefront subtracted from the estimated wavefront (created with the estimated modes)

3.2.2 Focal Plane Sharpening

Given the fact that FPS is not model dependent, (unlike phase diversity), it is worth investigating its performance under the same conditions as our phase diversity analysis. Our experience from our simulations suggests FPS can work with even lower SNR images than phase diversity. The obvious limiting factor will be the number of iterations (i.e. images) taken with FPS, which will be non-ideal if they are too many given the LWE varies over time. We therefore implement the realistic constraint that FPS should contain no more than 60 iterations (i.e. 1 total minute given each exposure is ∼1s on the DTTS) if it is to be considered reasonably for any real implementation. We hypothesize that a starting point with phase diversity may benefit from an improvement of FPS, based on its model independence. Furthermore, we hypothesize that the number of iterations (images) will be greatly reduced if using a starting point from phase diversity. Therefore we investigate 5 scenarios:

- Case1: Phase Diversity on a typical SNR DTTS image, taking the fastest solution - which is when the object is assumed known (i.e. scenario 3 from the previous section).
- Case2: Focal Plane Sharpening on the same type of image, starting from the null position.
- Case3: Focal Plane Sharpening performed on the solution from the Phase Diversity example.
- Case4: Focal Plane Sharpening performed on the best solution from the Phase Diversity example to explore if it does outperform the model dependent Phase Diversity
- Case5: Focal Plane Sharpening starting from the lowest SNR image possible

Table 1 summarizes the results from the aforementioned cases. It should be mentioned here that global tip and tilt are removed from the residual wavefront, where the residual is simply the estimated wavefront subtracted from the known aberrated wavefront. We find FPS improves the phase diversity result when we assume the object is known. However, we do not find any improvement from FPS on our best case phase diversity (where we used a longer exposure and estimated the object). In other words, in the case where the phase diversity images have a low SNR, FPS will always improve those images. We find the FPS improvement takes ~ 60 images when used with an initial phase diversity estimate, and for each image there is a very small computation time. Therefore we conclude this method would take between 1-2 minutes, which is slightly outside our prescribed time constraint of 1 minute.

Interestingly, we also find that FPS run by itself will converge on a solution much different than the original modal coefficients (see the rms residual coefficients column of Table 1), and furthermore this solution is of comparable residual WFE to our other best scenarios. The number of iterations for this convergence was about 120 - therefore double the images when using an initial starting point of phase diversity. Perhaps even more intriguing was the lower end SNR capabilities of FPS: we found FPS could successfully converge on images with a SNR of 10, in about 120 iterations.

3.2.3 Single image PD

Considering the evolution between two sequential images acquired in a closed loop AO system (due to evolving seeing and AO residuals), it is desirable to consider phase diversity using only 1 image. Work has been done on-sky in the past showing the problems two sequential phase diversity images can have on the algorithm.[10] Furthermore, the DTTS has a natural focus amplitude of 20 nm PV, providing a diverse image with no reference slope manipulation. We have found that when using a non-simple pupil - such as the case with the apodized VLT pupil considered here - that single image phase diversity works for any range of diverse values. This technique only works if the object is assumed to be known, in which case we assume a point source (which is not unreasonable considering we are simulating stars). In a companion SPIE paper (no. 9909-247), we consider the limitations of single image phase diversity (i.e. for uniform, circular, symmetric pupils). However, for this paper we will not explore the technical background of this technique. It is worth noting similar work has been done investigating single image Phase Diversity[11] in developing the LIFT technique, and it is not an entirely new concept.

Figure 4 summarizes our results: we see no major difference between 1 and 2 image phase diversity, as indicated by the similar rms coefficient residuals (top plot of figure). We also find that single image phase diversity does is not hindered by relatively low signal to noise ratios (around 70 in this case). Furthermore, it appears that even an in-focus image will work as the single image. Upon further investigation this ability relies entirely on the pupil, when we change to a circular pupil single image phase diversity breaks down for low amplitude diversities (see companion SPIE paper no. 9909-247). This falls in line with similar work,[12,13] where diversity has been implemented by manipulation of the pupil plane; in this case the pupil diversity originates from both the apodization and the non-perfect symmetry of the SPHERE pupil.

4. CASE STUDY 2: ESTIMATING THE SEGMENT PISTON ERRORS ON KECK

4.1 Basis and simulated images

Now we consider correcting segment-piston errors on Keck using the same approach in correcting the LWE in Section 3. We adopt a basis similar to the LWE basis, except in this case we consider piston only for each Keck segment (therefore a 36 element piston basis); the coefficients of this basis are to be estimated with Phase Diversity. The images are created using a 32x32 pixel pupil sampling. The phase estimation is virtually identical to the VLT scenario, with a few exceptions: we use phase diversity only (i.e. no focal plane sharpening), and we assume short-exposure 'frozen' turbulence images. Furthermore, we apply servo-lag, aliasing, dm-fitting, and photon effects to the phase of our AO-corrected image (the cumulative phase of these effects is shown in Figure 5). We again use OOMAO to generate our H-band images, which we try to create as similar to the NIRC2 detector as possible - this is the AO system we assume can make the diversity required for our images, and thus the system considered in this exercise. We assume the images are photon and read-noise dominated (adopting a realistic value of 60e for read-noise) and use a DM with 21 actuators across the pupil. The sampling of the resulting images is 3 pixels across the FWHM (close to the real NIRC2 sampling).

As previously discussed, Keck is known to have 'low order residual errors' which have been identified as piston phasing errors. We consider correcting a phase map representative of these errors, with a total WFE of ~ 150 nm rms. NIRC2 has a strong non-common path component of 60 nm rms WFE of 90 degree astigmatism, which we also add on top of the piston-segment errors. Both phase maps can be seen in Figure 5, along with the Keck pupil.

4.2 Estimation methods

For this example we employ single image phase diversity, with the assumption Keck has the capability of taking short-exposure images; the limitation of this assumption hinges on the detector capabilities as a wide selection of very bright stars are available on the sky. The motivation for single image phase diversity is for the same reasons described in Section 3.2.3. Furthermore, the linearity of the NIRC2 detector is limited to counts under $\sim$ 30000 ADU, which can be problematic when considering the two image phase diversity case, as the out-of-focus image will be entirely limited by the peak of the in focus image, which in turn is limited by the linearity of the

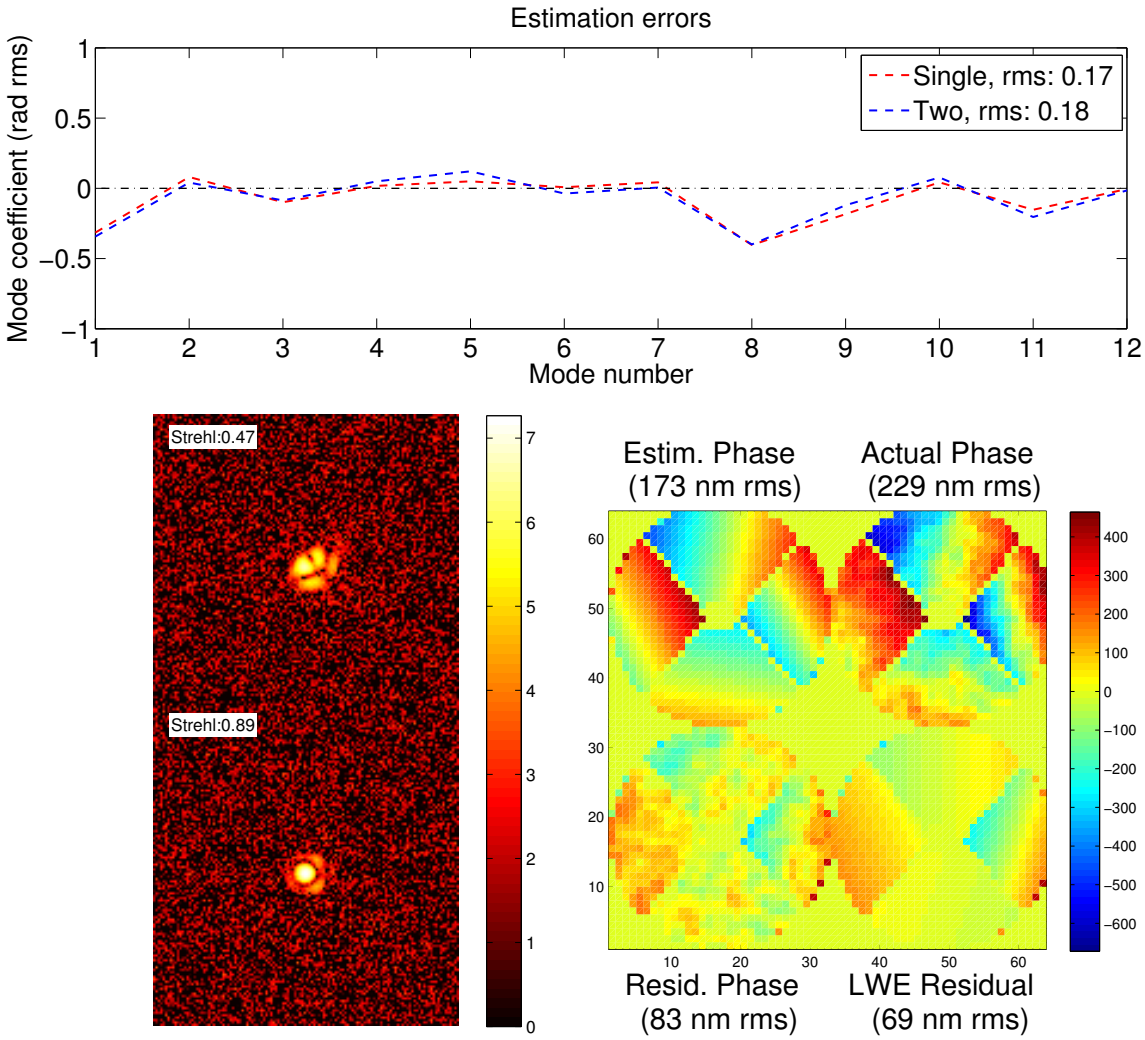

Figure 4. Top: residual coefficients from the actual coefficients for both single and two image phase diversity - both cases yield similar results. Bottom-left: original image and image after applying the phase correction estimated by single image phase diversity. Bottom-right: phase maps summarizing the residual of the estimation from single image phase diversity - the 'LWE Residual' map shows the result from subtracting the actual LWE phase from the estimated phase.

detector. There are ways to get around this, say by looking at a brighter star for the out of focus image, however we propose single image phase diversity to avoid this problem altogether.

An AO system has been simulated correcting a 21*21 actuator region across the telescope pupil, with realistic residual phase errors; these errors include DM lag, aliasing, photon noise, and fitting error. The single image used in our Phase Diversity is 1.5 waves of focus (which is $\sim$ 2.5 micron PV of stroke), which spreads the star's intensity to the edges of the controllable region of the DM; we use small field of view on our image, corresponding to only this correctable region. Finally, we ensure the ADU count of the diverse image is well under the 30000 limit (for the results of this work the maximum ADU count is $\sim$ 18,000).

The NIRC2 NCPA is around 60 nm rms of astigmatism; it is therefore convenient to try and estimate this error along with the segmented errors, with the goal of decoupling the two entirely such that segment correcting commands can be decoupled from DM NCPA commands. We note that the existing Phase Diversity algorithm at Keck is capable of estimating this astigmatism, however we chose to estimate here for an exercise. Figure 6 shows the results from single image phase diversity, simulated from an short exposure. The diverse image (shown on the middle left) has 1.5 waves of focus and falls well under the non-linear region of the NIRC2 detector (which

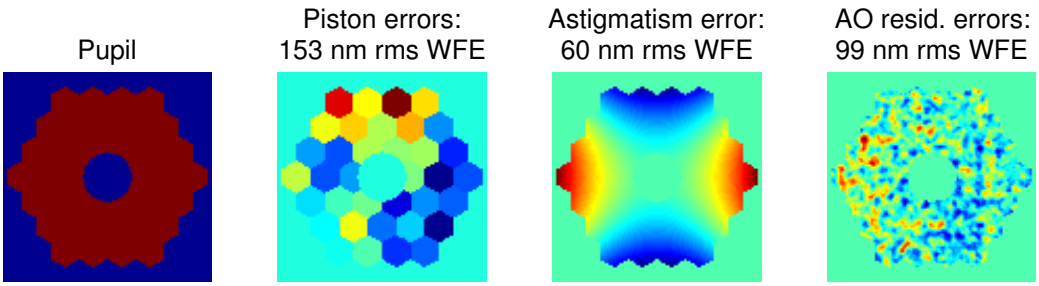

Figure 5. From left to right: Keck pupil, simulated piston phasing errors to be estimated, inherent astigmatism of NIRC2, AO phase errors (i.e. servo lag, aliasing, photon noise, fitting). The three phase maps on the right are used to create our simulated Keck images.

is around 30000 ADU). The 1.5 waves of focus was chosen because it yielded more accurate results than 1 wave, and for 2 waves the outer edge of the focus pattern crept into the uncorrected halo (see our companion SPIE paper no. 9909-247 for a description of how using 2 waves of focus yields better phase estimates). It can be seen from the figure that the estimated piston modes are very close the the actual piston errors, and furthermore that the astigmatism (Z_6) is recovered as well. The phase maps on the bottom of the Figure 6 show then when the actual piston error is subtracted from the estimate, the residual is almost exactly the 60 nm rms of astigmatism (shown by the bottom right phase map); the astigmatism is estimated to within 5% of the original value. It is left to determine if we could recover an NCPA profile estimate if the astigmatism were replaced by say 50 nm rms NCPA; we will be considering this in the future.

4.3 Other notes

Focal Plane Sharpening yields similar results to the Phase Diversity solution, although the evolution between sequential images is not properly modelled. We don't see this as a problem, considering FPS will work on extremely low SNR images as was shown in the LWE example (see Section 3.2.2), and as long as the FWHM of the star is not changing in the closed loop it should not be an issue. Furthermore, Focal Plane Sharpening takes around 120 iterations (therefore images) for convergence, and therefore might not be feasible for a real on-sky application.

Real images on Keck will have perhaps slightly longer integration times and it would useful to simulate images which average the AO residuals over that time. This will be tested in the near future.

Finally, single image phase diversity breaks down if too small of diversity is used here, whereas this was not the case with SPHERE. The only major difference between the two pupils in this example is the apodization, and the semi-symmetry of the VLT pupil. This lends insight on the capabilities of single image phase diversity, and how it might be useful with apodized/semi-symmetric pupils.

5. CONCLUSION

Segmented pupil error phase discontinuities were successfully estimated for two realistic scenarios on both VLT and Keck. From this work we conclude:

- For SPHERE: typical DTTS imaging conditions make it difficult for 'traditional' phase diversity, and we propose 3 solutions that work, when shown in simulation. The best correction was achieved by simply taking longer exposure images than typical DTTS signals, and estimating the object with two image phase diversity. However, this scenario did not model the long exposure of AO residual effects such as servo-lag and aliasing, and real images may be hindered by this. If this is the case then an approach that does not use as long of exposure times and still yields suitable LWE estimation is either by using a higher diversity mode, or leaving out the object estimation (by assuming a point-source). The latter method has an extremely fast convergence rate in comparison from the other methods.

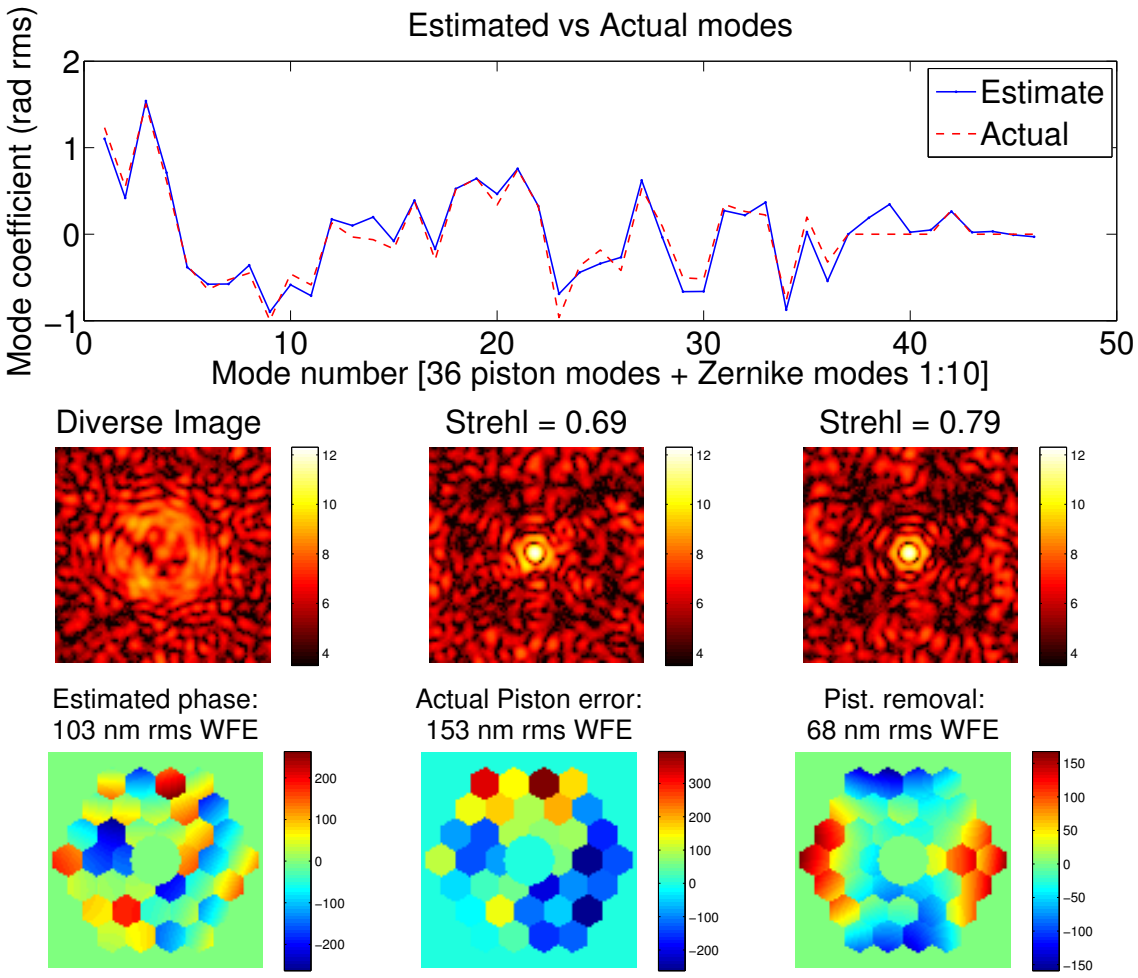

Figure 6. Correcting Keck segment errors with single image phase diversity. Top: Estimation of 36 piston modes in addition to the first 10 Zernike modes; it can be seen that not only are the segment modes estimated well but also the single astigmatism known in the system (we estimate residual tip/tilt but that does not effect the image). Middle from left to right: single diversity image (1.5 waves PV of focus, $\sim$ 10000 peak ADU, contained within the correctable region of the DM - this region is clearly seen in the right images), original image with all errors (piston, astigmatism, and AO residual errors), and the image after correcting for the estimate piston and zernike modes.

- If using the non-object estimation method described above, then Focal Plane Sharpening will improve the image (on top of the Phase Diversity estimation) in around 60 iterations, which is around 1-2 minutes assuming each image takes 1 second worth of time. This total time may be considered reasonable for a real implementation in correcting the LWE on sky. However, if starting from the null position, Focal Plane Sharpening can take more than twice this number of iterations. Worthy of note, however, is that Focal Plane Sharpening can work on extremely low SNR (around 10) DTTS images.
- Running low SNR FPS on top of low SNR phase diversity yields roughly the same estimate as the best case phase diversity scenario, where a high SNR (i.e. long exposure) is required.
- Single image phase diversity can correct the LWE to the same degree as two image phase diversity when not estimating the object. This could be useful when considering real on-sky images will have significant differences between eachother. Furthermore, we find that the diversity does not matter for the single image, and we believe this is attributed to the apodized pupil of SPHERE. We hope to investigate this last point in more detail in the future.

- If the LWE is estimated using a method such as ZELDA, it would be very useful to obtain the DTTS images and phase LWE estimation in order to check the validity of our methods.
- Phase segment piston errors were successfully estimated for Keck/NIRC2 simulated images given the assumption that images can be acquired near-instantaneous (i.e. with 'frozen' turbulence); this assumption should not be too unreasonable considering the limitation is only on the shutter speed of the detector (there should always be bright enough stars). The method used to estimate these errors was single image phase diversity. The inherent astigmatism was also simultaneously estimated with these segment errors and the astigmatism was recovered to within 5% of the original value. This is of particular importance when trying to de-couple segment errors from detector errors, where each will be corrected with different methods.

ACKNOWLEDGMENTS

The research leading to these results received the support of the A*MIDEX project (no. ANR-11-IDEX-0001-02) funded by the Investissements d'Avenir French Government program, managed by the French National Research Agency (ANR). This research was also funded in part by a MITACS/Campus France Globalink Research Award (ref: IT06712).

All the simulations and analysis done with the object-oriented MATLAB AO simulator (OOMAO) freely available from https://github.com/cmcorreia/LAM-Public

REFERENCES

[1] Rampy, R., Ragland, S., Wizinowich, P., and Campbell, R., "Understanding and correcting low order residual static aberrations in adaptive optics corrected images," in [*Adaptive Optics Systems IV*], **9148**, 91485I (Aug. 2014).

[2] Ragland, S., Jolissaint, L., Wizinowich, P., and Neyman, C., "Status of point spread function determination for Keck adaptive optics," in [*Adaptive Optics Systems IV*], **9148**, 91480S (July 2014).

[3] van Dam, M., "Effect of Phasing Errors on NIRC2 Images," in [*Keck Adaptive Optics Note 1117*], (Apr. 2016).

[4] N'Diaye, M., Dohlen, K., Fusco, T., and Paul, B., "Calibration of quasi-static aberrations in exoplanet direct-imaging instruments with a Zernike phase-mask sensor," **555**, A94 (July 2013).

[5] Paxman, R. G., Schulz, T. J., and Fienup, J. R., "Joint estimation of object and aberrations by using phase diversity," *Journal of the Optical Society of America A* **9**, 1027–1085 (1992).

[6] Lamb, M., Andersen, D. R., Véran, J.-P., Correia, C., Herriot, G., Rosensteiner, M., and Fiege, J., "Non-common path aberration corrections for current and future AO systems," in [*Adaptive Optics Systems IV*], **9148**, 914857 (July 2014).

[7] Conan, R. and Correia, C., "Object-oriented Matlab adaptive optics toolbox," in [*Adaptive Optics Systems IV*], **9148**, 91486C (Aug. 2014).

[8] Mugnier, L. M., Blanc, A., and Idier, J., "Phase diversity: a technique for wave-front sensing and for diffraction-limited imaging," in [*Advances in Imaging and Electron Physics*], Hawkes, P., ed., **141**, ch. 1, 1–76, Elsevier (2006).

[9] Lagarias., J., Reeds, J. A., Wright, M. H., and Wright, P. E., "Convergence Properties of the Nelder-Mead Simplex Method in Low Dimensions," *"SIAM Journal of Optimization"* **9**(1), 112–147 (1998).

[10] Jolissaint, L., Mugnier, L. M., Neyman, C., Christou, J., and Wizinowich, P., "Retrieving the telescope and instrument static wavefront aberration with a phase diversity procedure using on-sky adaptive optics corrected images," in [*Adaptive Optics Systems III*], **8447**, 844716 (July 2012).

[11] Meimon, S., Fusco, T., and Mugnier, L. M., "LIFT: a focal-plane wavefront sensor for real-time low-order sensing on faint sources," **35**(18), 3036–3038 (2010).

[12] Moore, D. B. and Fienup, J. R., "Sub-aperture position estimation in transverse-translation diversity wavefront sensing," in [*Imaging and Applied Optics 2015*], *Imaging and Applied Optics 2015* , AOM3F.4, Optical Society of America (2015).

[13] Brady, G. R., Guizar-Sicairos, M., and Fienup, J. R., "Optical wavefront measurement using phase retrieval with transverse translation diversity," *Opt. Express* **17**, 624–639 (Jan 2009).