

Journal of Astronomical Telescopes, Instruments, and Systems

AstronomicalTelescopes.SPIEDigitalLibrary.org

Quantifying telescope phase discontinuities external to adaptive optics systems by use of phase diversity and focal plane sharpening

Masen P. Lamb
Carlos Correia
Jean-François Sauvage
Jean-Pierre Véran
David R. Andersen
Arthur Vigan
Peter L. Wizinowich
Marcos A. van Dam
Laurent Mugnier
Charlotte Bond

Masen P. Lamb, Carlos Correia, Jean-François Sauvage, Jean-Pierre Véran, David R. Andersen, Arthur Vigan, Peter L. Wizinowich, Marcos A. van Dam, Laurent Mugnier, Charlotte Bond, "Quantifying telescope phase discontinuities external to adaptive optics systems by use of phase diversity and focal plane sharpening," *J. Astron. Telesc. Instrum. Syst.* **3**(3), 039001 (2017), doi: 10.1117/1.JATIS.3.3.039001.

SPIE.

Quantifying telescope phase discontinuities external to adaptive optics systems by use of phase diversity and focal plane sharpening

Masen P. Lamb,^{a,b,*} Carlos Correia,^c Jean-François Sauvage,^{c,d} Jean-Pierre Véran,^b David R. Andersen,^b Arthur Vigan,^c Peter L. Wizinowich,^e Marcos A. van Dam,^f Laurent Mugnier,^d and Charlotte Bond^c

^aUniversity of Victoria, Department of Physics and Astronomy, Victoria, Canada

^bNRC Herzberg Astronomy, Victoria, Canada

^cAix Marseille Université, CNRS, Laboratoire d'Astrophysique de Marseille (LAM) UMR 7326, Marseille, France

^dONERA, Optics Department, Châtillon, France

^eW. M. Keck Observatory, Kamuela, Hawaii, United States

^fFlat Wavefronts, Christchurch, New Zealand

Abstract. We propose and apply two methods to estimate pupil plane phase discontinuities for two realistic scenarios on the very large telescope (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) that is responsible for focal plane errors in the spectro-polarimetric high contrast exoplanet research (SPHERE) system in low wind and good seeing conditions. We successfully estimate the simulated LWE using both methods and show that they are complementary to one another. We also demonstrate that single image phase diversity (also known as phase retrieval with diversity) is also capable of estimating the simulated LWE when using the natural defocus on the SPHERE/differential tip tilt sensor (DTTS) imager. We demonstrate that phase diversity can estimate the LWE to within 30-nm root mean square wavefront error (RMS WFE), which is within the allowable tolerances to achieve a target SPHERE contrast of 10^{-6} . Finally, we simulate 153-nm RMS of piston errors on the mirror segments of Keck and produce NIRC2 images subject to these effects. We show that a single, diverse image with 1.5 waves (peak-to-valley) of focus can be used to estimate this error to within 29-nm RMS WFE, and a perfect correction of our estimation would increase the Strehl ratio of an NIRC2 image by 12%. © 2017 Society of Photo-Optical Instrumentation Engineers (SPIE) [DOI: [10.1117/1.JATIS.3.3.039001](https://doi.org/10.1117/1.JATIS.3.3.039001)]

Keywords: phase diversity; image sharpening; wavefront sensing; adaptive optics.

Paper 16056P received Nov. 1, 2016; accepted for publication May 24, 2017; published online Jun. 16, 2017.

1 Introduction

Piston discontinuities in segmented pupils are difficult to quantify when considering traditional adaptive optics (AO) systems, which originate from the inability of the Shack–Hartmann wavefront sensor (SHWFS) to estimate differential piston. A prime example of this is the low wind effect (LWE) on the very large telescope (VLT)/spectro-polarimetric high contrast exoplanet research (SPHERE) system, where nights with good seeing and a relative lack of wind surprisingly yield focal plane images of poor quality. This effect has been interpreted to be a result of temperature discontinuities across the VLT pupil.^{1,2} Conceptually, these temperature discontinuities are thought to be defined by the secondary mirror spiders, which act as thermal barriers between segments. Airflow simulations of the spiders have been shown to reproduce these effects, and for more information, we refer the reader to Sauvage et al.² Sharp differentials in temperature may translate to a differential piston effect, which is due to the index of refraction of air having a dependency on temperature. The SHWFS is simply unable to detect differential piston and the resulting point spread function (PSF) reveals features akin to “Mickey Mouse ears”¹ (see Fig. 1); correcting this effect is crucial considering the purpose of SPHERE is to achieve the highest contrast possible, which is clearly contaminated by this effect. The target contrast of

SPHERE is 10^{-6} at $0.5''$,³ which is achievable if the noncommon path aberrations (NCPA) in the system are under 50-nm RMS WFE (wavefront error).² The internal NCPA of SPHERE was initially quantified at 25-nm RMS;⁴ however, this value is thought to have grown to ~ 40 -nm RMS and will be quantified in the near future. Assuming this latter value to be true, this leaves a maximum of only 30-nm RMS (the quadratic sum) from other contributions such as the LWE to achieve SPHERE’s target contrast. The LWE is observed to occur when nights have submeter per second speeds ($\sim 1/5$ of the nights at Cerro Paranal), therefore, such a method to quantify and correct this effect could be very valuable for the proper functionality of SPHERE. This effect drives the operation of the observatory to disable the SPHERE instrument and reorganize the observations program. Developing the ability to estimate the LWE to within 30-nm RMS is, therefore, critical for the success of the SPHERE project.

Piston discontinuities also impact the performance of segmented telescopes, such as the Keck telescopes. Artifacts of the PSF due to differential piston (and potentially other sources) on the Keck/NIRC2 AO system have been observed and identified as “low-order residual errors”^{5,6} (see Fig. 1). These errors typically result in a PSF with a deviation in the first diffraction ring and a reduced Strehl ratio and are shown to be stable

*Address all correspondence to: Masen P. Lamb, E-mail: masen@uvic.ca




Fig. 1 (a) An image acquired by the DTTS imager on SPHERE during a night with a strong LWE, shown in log-scale (courtesy of Sauvage). The asymmetric “ear”-like features on the PSF shown here are an example of the PSF contamination experienced throughout the course of the entire night and restrict the use of the instrument. (b) A K-band, short exposure Keck/NIRC2 image (also shown in log-scale) displaying typical features of “low order residuals” that are persistent throughout the duration of the closed AO loop (courtesy of Ragland).

on relatively long time scales (i.e., >30 min⁷). Some, but not all,⁷ of these errors arise from the inability of the SHWFS to detect the differential piston. Any correction of this effect requires a reliable estimation, applying such a correction will increase the Strehl ratio and allow the AO system to achieve its full potential. Estimating this effect will also provide essential information for PSF reconstruction, which is vital for a variety of astronomical applications.⁸

These two pupil effects have been investigated using a variety of different approaches. For example, recent studies (see N’Diaye et al.⁹) have explored the use of a Zernike WFS (ZELDA¹⁰), capable of detecting these piston variations on VLT/SPHERE. In the case of the Keck low-order residuals, the approaches of both the Gerchberg–Saxton algorithm and phase diversity have been used to estimate these residuals both in simulation and on-sky.^{7,8}

In this paper, we explore two different methods for estimating these pupil-discontinuity effects for simulated on-sky data. The first method is the well-established approach of phase diversity, where known diverse images are compared with their synthetic counterpart to estimate the phase of an optical system.¹¹ The second method is the technique of focal plane sharpening¹² (FPS), where the PSF in the focal plane is optimized using only 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 to estimate the respective errors in question. We also explore the concept of single image phase diversity (also known as phase diverse phase retrieval), which could be very useful in avoiding turbulence evolution and AO residuals between a set of on-sky images.

2 Estimation Methods

2.1 Phase Diversity

We employ a phase diversity code that follows the formulation of Paxman et al.,¹³ where an aberration-only objective function and its gradient are fed through a nonlinear optimization algorithm to minimize the quadratic difference between synthetic and observed 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 is described and

determined in Lamb et al.;¹⁴ when the quadratic difference reached between the images is below this value convergence. The synthetic and real data are typically in and out of focus images, although the code can incorporate any number of images or type of diversity. The optimization technique we employ is the well-established quasi-Newton Broyden–Fletcher–Goldfarb–Shanno 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). This code has been developed as a class for the AO MATLAB software object-oriented MATLAB AO (OOMAO) simulator;¹⁵ for more details of the code, see Lamb et al.,¹⁴ and for an overview of the phase diversity technique, we refer the reader to Mugnier et al.¹⁶ 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

The results in this work use a MATLAB-based focal plane sharpening code, which has also been developed as a class for OOMAO¹⁵ and has cross-compatibility with its phase diversity counterpart. 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;¹⁷ 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). This method, along with a description of different criteria choices, is explained in more detail in Lamb et al.¹⁴ The results in this work use the following criteria: at each step in the optimization, a small region centered on the PSF is extracted and subsequently median-filtered with a $2 \times$ FWHM kernel (i.e., two times the number of pixels across the FWHM of the theoretical diffraction limited PSF) to reduce noise; a 2-D Gaussian is fit to this image from which the amplitude is measured. The magnitude of this amplitude is the metric that is optimized, changing at each iteration with the new set of basis coefficients.

3 Estimating the Low Wind Effect on Sphere

3.1 Basis and Simulated Images

To estimate the LWE on SPHERE, we propose a basis roughly defining the pupil plane phase variations that occur at each quadrant of the VLT pupil (see Fig. 2). 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 the performance of each method. The amplitude of the piston errors we use to artificially produce the LWE is ~ 1200 nm peak-to-valley (PV) WFE and is shown in Fig. 3 (right); this LWE error is conservatively high, as typical LWE errors exist in the range of 600 to 800-nm PV WFE.¹ The objective of our work is to estimate this error to within 30-nm RMS. We also adopt a typical NCPA representative of the SPHERE system, specifically 45-nm RMS WFE following a $1/\nu^2$ power law; they are also shown in the same figure. As described in Sec. 1, the magnitude of the SPHERE NCPA is estimated to be ~ 40 -nm RMS, and our choice of 45-nm RMS is chosen to reflect a conservative overestimate of the error. This overestimate increases the error budget to 54-nm RMS, and we still aim to quantify the LWE to within 30-nm RMS WFE.




Fig. 2 Piston, tip, and tilt basis used to recreate the PSF variations seen during the LWE on SPHERE. Each mode is normalized to 1 rad RMS (except the pistons). For the remainder of this paper, mode “1” of this basis corresponds to the top left mode shown here (piston on the left segment). The remaining modes numerically follow from left to right, ending with mode “12” shown in the bottom right of this figure (tip on the top segment).




Fig. 3 (a) VLT pupil, (b) SPHERE apodization mask, (c) assumed NCPA corresponding to 45-nm RMS WFE, and (d) 1200-nm PV WFE LWE errors.

Phase diversity and FPS require focal plane images of the PSF in order to estimate the coefficients of this selected basis; these images are created as simulations of the SPHERE differential tip tilt sensor (DTTS) imager. The DTTS imager exists directly before the coronagraph on SPHERE and thus shares most of the common internal aberrations with the science image optics. Furthermore, using the DTTS images for the LWE characterization is much less complicated than using the coronagraphic images, even though the latter can be used for the estimation of aberrations.^{18,19} We, therefore, aim to show the DTTS imager is sufficient for the LWE estimation and create synthetic DTTS H-band images dominated by photon and read noise.

The images are created using the AO simulation tool OOMAO with a 32×32 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 (~ 4 pixels across the FWHM). We subject the image to a turbulence profile (generated assuming an r_0 of 11 cm at 550 nm, and that all the turbulence occurs at the ground layer), and subsequently apply an AO correction using the simulated DM. The turbulence is evolved at a sampling rate of 500 Hz and with a wind speed of 15 m/s; long exposure images are created by stepping the turbulence over the appropriate number of sampling steps pertaining to the total integration time of the image. This is particularly important as the DTTS imager typically acquires ~ 1 s exposures. However, we note 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 FPS 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 have found 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” correctable region.

3.2 Estimation Methods

3.2.1 Phase diversity

We first consider the estimation of the LWE modal coefficients by means of phase diversity; in particular, we employ what we consider “classic” phase diversity—when two images are used with focus diversity and the object is simultaneously estimated. Due to the combination of 45-nm RMS NCPA and the relatively large PV amplitude of the LWE (~ 1200 PV nm), we consider two waves PV of focus to ensure the diversity is larger than the phase to be estimated.

It is important to simulate images with realistic DTTS signal-to-noise ratios (SNR), thus we consider the on-sky data shown in Fig. 4. We simulate a star with a typical SNR from this plot and run phase diversity with the aforementioned parameters. It was found that we have difficulty on convergence, where half




Fig. 4 DTTS imager data, taken from SPHERE. The different colors correspond to different acquisition modes: blue points are taken in a mode optimized for bright stars, whereas red points are suited for fainter stars; the green points represent an additional acquisition mode rarely used (and therefore explains the lack of points in this plot). The cloud of points around 2 to 3 ADU correspond to misdetections, 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 an SNR ~ 70 and use this value for our analysis.

the time the solution will “run away” in tip and tilt and converge on an erroneous result. We consider three approaches to solving this issue:

1. Increase the SNR of the image (i.e., increasing the exposure time of the image), however, on the real system, this will integrate residual AO effects such as lag error, etc. We find this solution increases the Strehl from 0.47 to 0.96 (see Fig. 5), 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 option.
2. The convergence seems to break down because of a signal-to-noise issue; therefore, it is worth exploring how a different type of diverse image performs under this same noise. Higher diversity modes, specifically cousins of the trefoil family (i.e., Z_{11} , Z_{21} , etc.), seem to work better in general in our simulation. In particular, we consider Z_{66} , where we have chosen this mode because it is a relatively high-order mode—which we find in general produces better phase estimates, presumably because of the large diversity introduced to the PSF—and it is not too high such that it will be difficult to create with a SPHERE-like DM. Our simulation shows this mode always produces a better estimate of the phase (by a few nm RMS), and it converges faster than the focus-diversity case. We adopt two waves (PV) of Z_{66} as our diversity, and our simulation shows a Strehl increase from 0.47 to 0.91 [see Fig. 5 (bottom left image)]. This improvement is not as high as the previous example; however, it is achieved with a lower SNR. One problem may arise

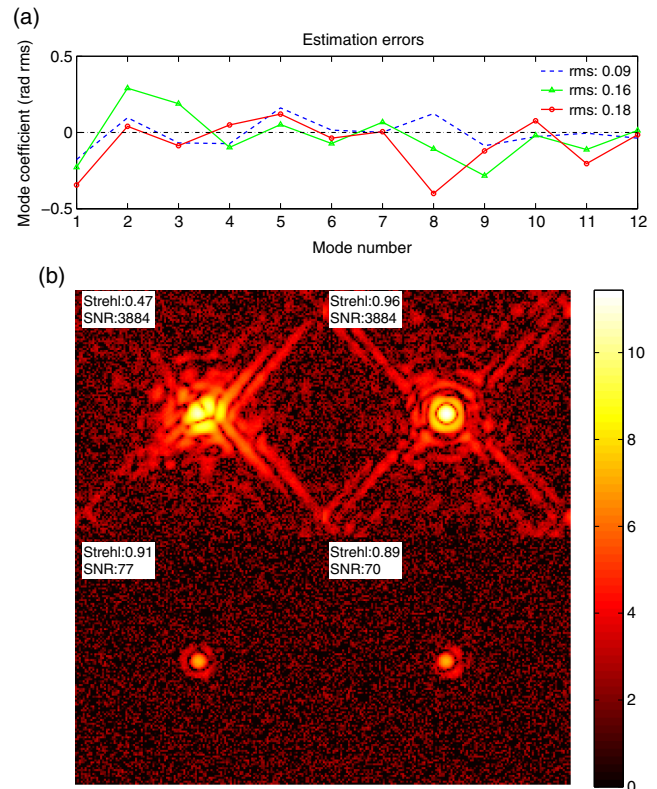


Fig. 5 (a) residuals of LWE piston, tip, and tilt estimations from the actual modes, using phase diversity for three different scenarios (blue dashed line: long exposure object estimation with focus diversity, green with triangles: object estimation with higher diversity (Z_{66}), and red with open circles: assumed point source with focus diversity). It can be seen that the long exposure (blue dashed line) scenario performs the best as indicated its RMS residual from the actual modes. (b) Four panels of simulated VLT images, created from the phase projection of the estimated modes; they are described as follows: the upper two panels include no correction and long exposure phase diversity, respectively. The lower two panels include the higher diversity and assumed point source scenarios, respectively. 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.

in realistically implementing this approach; high-order phase speckles introduced from this diversity 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 fitting error compared with a simpler mode, such as focus.

3. Assume that 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 (achieving a lower Strehl of 0.89 compared with the two aforementioned solutions); however, it seems to have a much faster convergence rate. Figure 5 shows results from our simulation using this method.

From all three of these scenarios, we conclude:

- Under typical DTTS imaging conditions, “classic” 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 by assuming a known object. In the former case, the estimate is more accurate, and in the latter, the speed is considerably faster.

It is worth noting that the diversity values adopted in this work were mainly chosen due to convention. In the context of phase diversity and phase retrieval, the exploration of induced diversity has been investigated by numerous examples in the literature. For example, Jurling and Fienup²⁰ consider scenarios with ± 8 waves of focus diversity. Therefore, we acknowledge that it is possible alternative diversities may be more desirable, but for the purposes of this work, we do not further pursue this topic.

3.2.2 Focal plane sharpening

Given the fact that FPS is not model dependent (unlike our form of phase diversity), it is worth investigating its performance under the same conditions as our phase diversity analysis. In general, our simulations suggest FPS works with even lower SNR images than phase diversity. However, one obvious limiting factor of FPS is the number of iterations (i.e., images) taken, particularly with the LWE considering the effect can vary over short time scales (i.e., tens of minutes). The LWE is observed to have time-varying evolution over the course of several minutes (see Sauvage et al.¹), and we arbitrarily choose a reasonable “window” in which to quantify this effect as 1 min. Therefore, we implement the constraint that FPS should contain no more than 60 iterations, given the fact that a typical DTTS exposure is ~ 1 s. In addition, we hypothesize that starting with an initial estimate from 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 five scenarios:

- Case 1: 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).

- Case 2: Focal plane sharpening on the same type of image, starting from the null position.
- Case 3: Focal plane sharpening performed on the solution from the phase diversity example.
- Case 4: Focal plane sharpening performed on the best solution from the phase diversity example to explore if it does indeed outperform the model-dependent phase diversity.
- Case 5: Focal plane sharpening starting from the lowest SNR image possible.

Table 1 summarizes the results from the aforementioned cases. We note here that global tip and tilt are removed from the residual wave front, where the residual wave front is the known aberrated wave front (LWE + NCPA) subtracted from that of the estimated wave front. These residual wave fronts include the 45-nm RMS NCPA and will, therefore, be much larger than our target LWE estimation of 30-nm RMS. In cases 4 and 5, FPS uses the PSF created from correcting the phase diversity estimate as input, in addition to using the phase diversity estimated modes as a starting position. It is found that FPS improves the phase diversity result when the object is assumed known. However, there is no clear improvement from FPS on the best case phase diversity from Sec. 3.2.1 (where we used a longer exposure and estimated the object). In other words, in the case where the phase diversity images have a lower SNR, FPS will always improve those images. We find that 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 that this method would take between 1 and 2 min, which is slightly outside our prescribed time constraint of 1 min.

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—about double the number of images that used an initial estimate from phase diversity. Perhaps even more intriguing was the lower end SNR capabilities of FPS: it

Table 1 Phase diversity and focal plane sharpening results correcting for the LWE.

Case	Method	Residual coefficients ^a RMS (rad)	Residual wave front ^b RMS (nm)	Strehl (Marechal)	No. of images	Initial SNR
Case 1	PD-1 (assumed object)	0.17	89.4	89.1	2	70
	PD-2 (long exposure) ^c	0.09	55.3	95.7	2	3.8e4
Case 2	FPS	0.40	56.1	95.5	122	106
Case 3	FPS + PD-1	0.12	57.9	95.3	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	92.1	121	10

^aThe RMS of the difference between the estimated and actual coefficients.

^bThe total aberrated wave front (LWE + NCPA) subtracted from the LWE estimate.

^cShown for reference.

was found that FPS could successfully converge on images with an SNR of 10, in about 120 iterations.

Regardless, we find that the best estimate of the LWE requires phase diversity with object estimation. For the remainder of Sec. 3, we further expand this method to assess its feasibility to accommodate SPHERE's WFE requirements as outlined in Sec. 1. In addition, we consider an alternative phase diversity scenario that requires only the raw DTTS images.

3.2.3 Estimating the LWE from a single image

Considering the evolution between two sequential images acquired in a closed-loop AO system (due to changing seeing and AO residuals), it is desirable to consider phase diversity using only one image. Single-image phase diversity is analogous to phase retrieval,¹¹ and the image in question is subject to diversity; for the rest of this paper, we will refer to this technique as “single-image phase diversity” instead of phase retrieval to stay consistent with nomenclature of the multiple image scenarios. Work has been done on-sky in the past, showing the challenges involved with using two sequential phase diverse images.²¹ Furthermore, the DTTS has a natural focus amplitude of 20-nm RMS, providing a diverse image with no reference slope manipulation. We have found that when using a nonsimple pupil—such as the case with the apodized VLT pupil considered here—single image phase diversity works for this amplitude of natural focus. This approach assumes the object is known, in which case we assume a point source (which is not unreasonable considering we are simulating stars). In Lamb et al.,¹⁴ we consider the limitations of single-image phase diversity (i.e., for uniform, circular, and symmetric pupils). However, for this paper, we will not explore the technical background of this technique. It is worth noting that a similar work has been done for investigating the single image phase diversity²² in developing the LIFT technique, and it is not an entirely new concept. Furthermore, the concept of a nonsimple pupil to facilitate phase estimation dates as far back as 1965 by Mehta.²³

To validate the performance of single-image phase diversity, we consider three different scenarios of LWE estimation

- i. “Classic phase diversity,” in which case phase and object are simultaneously estimated from two images with 0 and +2 waves PV of focus.
- ii. Two image phase diversity with no object estimation, using images with -1 and $+2$ waves PV of focus.
- iii. One image phase diversity, using a single-image subject to 20-nm PV focus, similar to the actual DTTS images.

The images used in the above scenarios have an SNR of ~ 70 for an in-focus image; this is to stay consistent with Sec. 3.2.1. In the case of classic phase diversity, however, the solution did not converge with this value and a minimum SNR of ~ 700 was required. Figure 6 summarizes the results of each of these scenarios. The best estimation is achieved by classic phase diversity; the residual between the phase estimate of the LWE and the actual injected LWE is 30-nm RMS (top right of the figure). However, in the cases of (ii) and (iii), LWE estimations are marginally worse with 50- and 62-nm RMS residuals (respectively). Two-image phase diversity with no object estimation is considered here strictly as a comparison with single-image phase diversity, where the only real difference between the two scenarios is

a single image. From these results, it appears single-image phase diversity with the natural DTTS focus can reasonably estimate the LWE, although not at the same performance of classic phase diversity.

3.3 Performance Evaluation

Through consideration of both FPS and phase diversity, it appears the most viable form of LWE estimation (in terms of both speed and accuracy) is in some form of phase diversity. As mentioned in Sec. 1, 30-nm RMS WFE of the LWE is small enough to still achieve the target contrast of SPHERE, therefore, a perfect correction of our estimate from a classic phase diversity would achieve this requirement. We note here that the LWE in this simulation is representative of a larger-than-normal LWE night; therefore, we expect better performance under less severe conditions. That being said, classic phase diversity requires DTTS images with a higher SNR than those delivered by the operation of the sensor and thus would be subject to atmospheric residuals over the course of long exposures. Furthermore, the acquisition of diverse images would require reference slope manipulation; it is worth considering that single-image phase diversity requires no such manipulation when using the natural DTTS focus. Furthermore, the results shown here suggest that the impact in LWE estimation is not terribly drastic between single image and classic phase diversity (i.e., on the order of 5% Strehl loss using the Marechal approximation) and could be seriously considered for estimation/correction if the science does not require the full 10^{-6} contrast. We are currently analyzing a sequence of DTTS images recently acquired at the VLT during a night subject to a strong LWE in order to better understand the feasibility of single-image phase diversity. The results of this study will be published in the near future.

4 Estimating the Segment Piston Errors on Keck

As previously mentioned, significant low-order AO residuals in the Keck/NIRC2 system exist and may originate in the form of cophasing errors of the primary mirror segments. We now consider estimating these segment-piston errors employing the same single-image phase diversity approach that was used in quantifying the LWE in Sec. 3.3. The motivation for this approach is the desire to avoid unwanted evolutionary effects between two sequential images, as has been discovered with classic phase diversity in the past (as previously mentioned in Sec. 3.2.3). We note that previous work has aimed at mitigating these effects by considering long exposure images with phase diversity, which average out the AO residuals and seeing effects;²⁴ however, in this work, we are interested in assessing the feasibility of using short exposure images. As we are interested in avoiding evolutionary effects between sequential images, we do not consider focal plane sharpening in this exercise.

A defocused NIRC2 image, subject to 153-nm RMS WFE cophasing errors, is simulated as faithfully as possible. The cophasing error injection we adopt here was taken from a similar phasing-residual analysis (using different estimation algorithms);⁷ the cophasing error phase map is shown in Fig. 7. The simulated images were created on a pupil with 32×32 pixel sampling and were subject to photon and read-noise errors; we consider a read-noise of $60e^-$ (a single read-out NIRC2 image is closer to $40e^-$ (Ref. 25), however, we choose a larger value to be conservative). The images are generated with

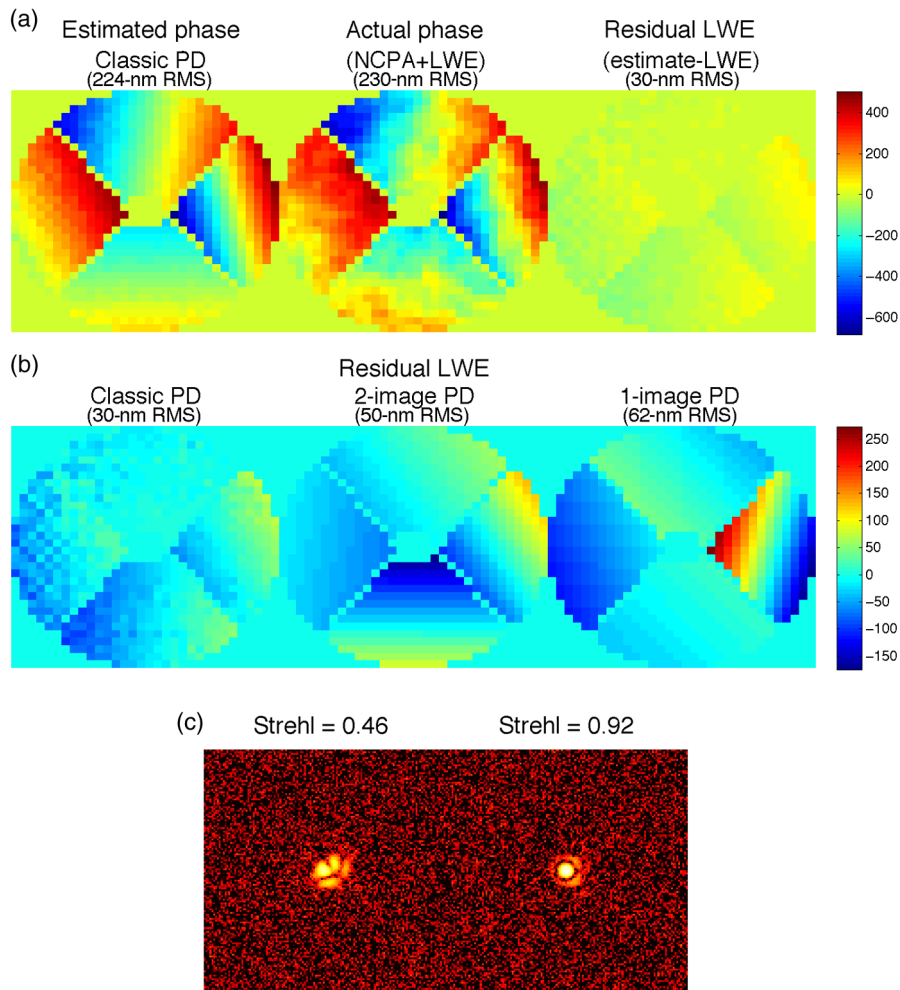


Fig. 6 (a) Estimated LWE (left) from classic phase diversity (phase and object estimation using images with 0 and +2 waves PV focus), actual phase injected (center), and residual phase between the estimate and the actual LWE injection (right). The residual WFE reaches the desired 30-nm RMS, such that a perfect correction of this estimated phase would result in a contrast at least 10^{-6} . (b) Residual phase maps for two-image phase diversity with and without object estimation (left and center panels, respectively) and single-image phase diversity (right panel, using a single image with the natural focus of the DTTS imager). These additional phase diversity scenarios do not meet the performance of classic phase diversity but are shown here for comparison. The case of the single image should be considered useful for its potential of both a quick LWE quantification and unobtrusiveness in image acquisition. (c) Simulated PSFs before and after (perfect) correction from the single-image LWE estimate.

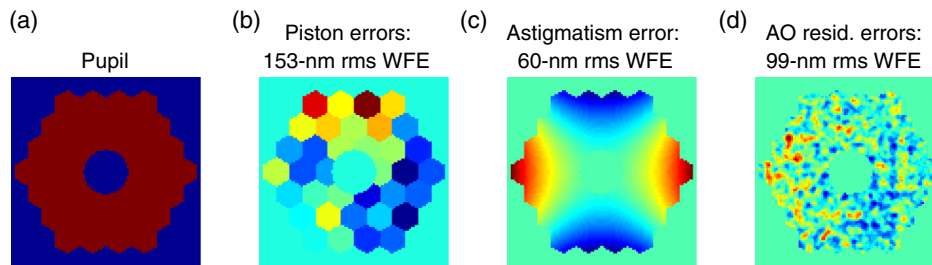


Fig. 7 (a) Keck pupil, (b) simulated piston phasing errors to be estimated, (c) inherent astigmatism of NIRC2, and (d) AO phase errors (i.e., servo lag, aliasing, photon noise, and fitting). The three phase maps on the right are used to create our simulated Keck images.

OOMAO in the environment of a simulated AO system with 21 actuators across the pupil (identical to the actual NIRC2 system). In addition to the cophasing residuals, we also inject 60-nm RMS of astigmatism (a rough representation of the true

noncommon path error) and 99-nm RMS of simulated AO residual errors (summarizing contributions from servo-lag, aliasing, and DM-fitting errors). Figure 7 shows phase maps representing these errors.

Assuming the object is a point source, the de-focused NIRC2 image is estimated with our phase diversity algorithm. The amount of defocus is 1.5 waves (PV) and was chosen because any larger values caused a focus pattern that crept into the uncorrected halo. The estimated basis modes are pure piston variations, defined by each mirror segment (therefore, a 36 element basis), along with the first 10 Zernike modes. One of the goals of our estimation is to see if it can decouple the piston errors from the astigmatism while under the influence of realistic AO residuals. Figure 8 summarizes our findings. The top of the figure shows the coefficients of the estimated piston plus Zernike basis in blue, overlaid on the actual modes injected into the system (red). The RMS residual difference between the actual and the estimated modes is 0.1092 rad,

which translates to 29-nm RMS. To visualize this estimation in terms of phase errors, the middle section of the figure displays the estimated phase (left), the actual NCPA plus segment error injection (middle), and residual phase of 29-nm RMS (right). Here the global tip and tilt is removed from both the estimated and the actual phase prior to the subtraction. Hence, the magnitude of the original 153-nm RMS cophasing error has been reduced to 117-nm RMS; when this error is added to the 60-nm RMS of NCPA, the final WFE is 99-nm RMS.

4.1 Performance Evaluation

To get a sense of the performance of this analysis, we consider the situation where a perfect correction of the phase estimation

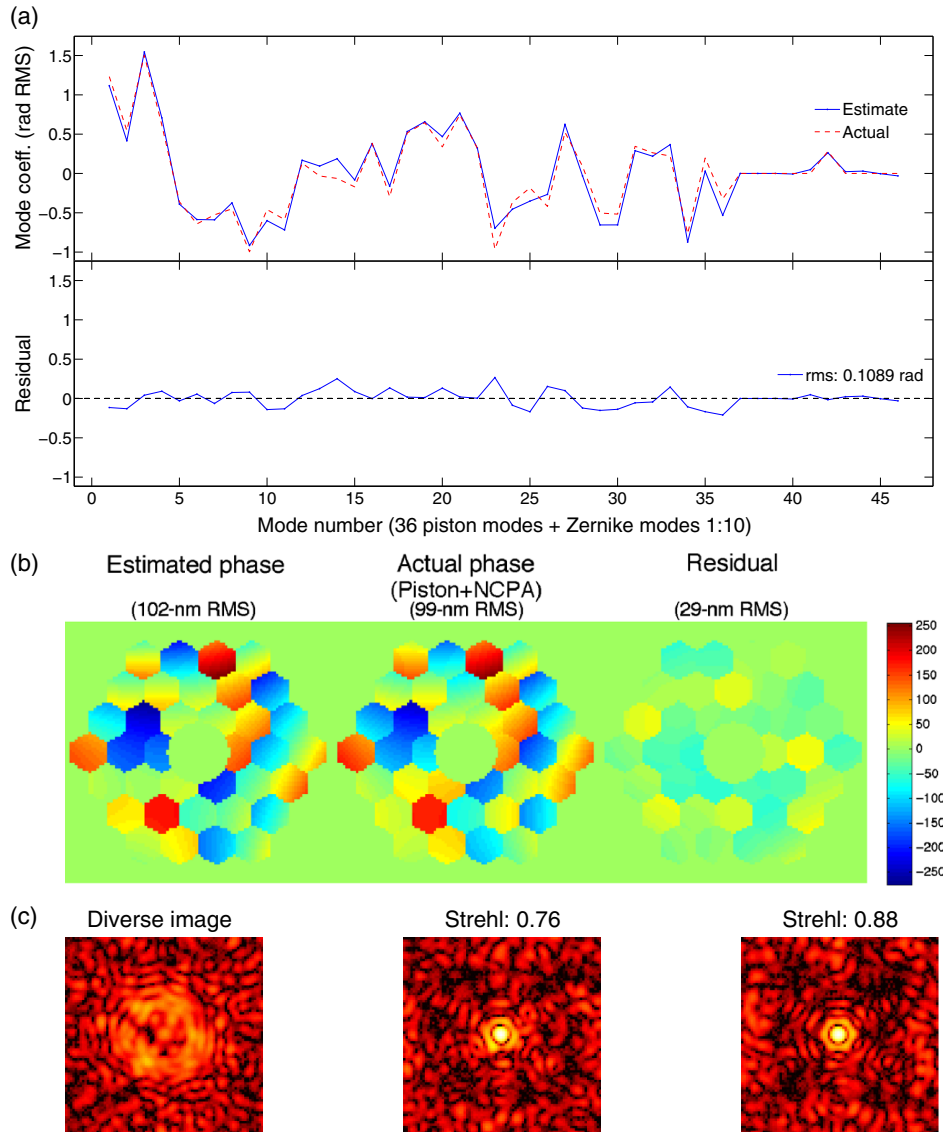


Fig. 8 Estimation of both segment piston errors and NCPA (astigmatism) from a single diverse image of a simulated bright star (diverse image shown in bottom left). (a) Estimated modal coefficients of 36 piston modes and 10 Zernike modes (Z_1 to Z_{10}). (b) Phase reconstructed from the estimated modes (left) compared with the actual piston plus astigmatism phase (middle); the residual between the two is shown on the right with a WFE of 29-nm RMS. The global tip and tilt was removed from the estimated and actual phases, reducing the original cophasing error from 153- to 117-nm RMS. (c) Simulated images of the initially aberrated system (middle) and situation where the perfect correction of the estimated phase is achieved (right).

can be achieved. The bottom of Fig. 8 displays an image before and after such a correction, showing a Strehl improvement of $\sim 12\%$ (calculated directly from the images, using the ratio between the image optical transfer functions and their diffraction limited counterpart). The signal to noise of a faint source increases proportionally to the Strehl ratio; therefore, the exposure time is proportional to the square of the Strehl. In our example, assuming a perfect correction of the segment-piston estimation, the exposure time of a faint object would be ~ 0.75 times less to achieve the same signal to noise without correction.

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 classic phase diversity, and we propose three solutions that improve the estimation, which we show in simulation. The best estimation was achieved by taking longer exposure images than typical of the DTTS, in addition to object estimation. However, our simulation does not model the long exposure AO residual effects such as servo-lag, aliasing, and fitting error, which would provide additional uncertainties in a true estimation. If, in reality, these residuals have a larger than desired impact on the estimation, we propose a workaround where a suitable LWE estimation is achieved by either using a higher diversity mode or leaving out the object estimation (and assuming a point-source). The latter method is less computationally expensive with respect to any phase plus object estimation.
- If using the nonobject 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 to 2 min assuming each image takes ~ 1 s. 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 (~ 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 estimation, where a high SNR (i.e., long exposure) is required.
- We find that single-image phase diversity is a useful tool to estimate the LWE. If using a single image with the natural focus on the DTTS imager, we find that this technique can estimate the LWE to almost the same degree as two images with no object estimation (using $+1$ and -2 waves of focus). However, two images with object estimation significantly improve the phase estimation, although higher SNR on the images is required. The single-image approach could be useful considering sequential on-sky images will have evolution from one image to the next—an effect known to cause issues with phase diversity in the past.²¹ Lamb et al. (in preparation) are currently analyzing on-sky DTTS data taken during an LWE sequence to further explore the feasibility of single-image phase diversity.
- To reach the target contrast of 10^{-6} , the LWE must be estimated to within an error of 30-nm RMS as discussed in Sec. 1. We find that this goal can be achieved when using classic phase diversity (an in/out of focus image with object estimation), given the SNR of the in-focus image is ~ 700 . The error on this estimation roughly doubles when we use phase diversity with a single DTTS image with a natural 20-nm RMS of focus, although on an image with an SNR 70. We conclude that at a small sacrifice in Strehl ($\sim 4\%$) and therefore contrast, single DTTS images with no manipulation could be used to quantify the LWE.
- The correction of the LWE effect using the wavefront control capabilities of SPHERE is a complex point and is not studied in this paper. The SPHERE system will benefit from a change of the spider coating, which will largely decrease the amplitude of the effect. Ideally, the addition of this coating, combined with some sort of wavefront control scheme from the LWE estimation will be able to correct for the total effect in order to achieve the design contrast of SPHERE. Some recent tests have been investigated by Sauvage to prove the SPHERE wavefront control is able to produce the corresponding amplitude estimates of the LWE, and the results from this study will be published in the near future.
- 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 (the availability of bright stars should not be an issue). A goal of this exercise was to show the capability of using phase diversity with a single, diverse image; this technique avoids any significant evolution effects between sequential images and would be a useful comparison to the complementary approach of using long exposure images for phase estimation.¹⁶ Given an aberrated image subject to 153-nm RMS WFE of simulated cophasing errors, we find that we can estimate the phase to within an error of 29-nm RMS. Assuming a perfect application of this estimate can be applied to the mirror segments, this would result in a Strehl increase of 12% for our simulated NIRC2 images. A direct result of such a fix would decrease the exposure time on faint sources by a factor of 1.3 to achieve the same SNR if no fix was applied.

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 freely available from <https://github.com/cmcorreia/LAM-Public>.

References

1. J.-F. Sauvage, “Low wind effect, the main limitation of the sphere instrument,” in *Adaptive Optics for Extremely Large Telescopes 4—Conf. Proc.*, Lake Arrowhead (2015).

2. J. F. Sauvage et al., "Tackling down the low wind effect on SPHERE instrument," *Proc. SPIE* **9909**, 990916 (2016).
3. J. L. Beuzit et al., "SPHERE: a 'Planet Finder' instrument for the VLT," *Proc. SPIE* **7014**, 701418 (2008).
4. J. F. Sauvage et al., "SAXO: the extreme adaptive optics system of SPHERE (I) system overview and global laboratory performance," *J. Astron. Telesc. Instrum. Syst.* **2**, 025003 (2016).
5. R. Rampy et al., "Understanding and correcting low order residual static aberrations in adaptive optics corrected images," *Proc. SPIE* **9148**, 91485I (2014).
6. S. Ragland et al., "Status of point spread function determination for Keck adaptive optics," *Proc. SPIE* **9148**, 91480S (2014).
7. M. van Dam, "Effect of phasing errors on NIRC2 images," in *Keck Adaptive Optics Note 1117* (2016).
8. S. Ragland et al., "Point spread function determination for Keck adaptive optics," *Proc. SPIE* **9909**, 99091P (2016).
9. M. N'Diaye et al., "Calibration of quasi-static aberrations in exoplanet direct-imaging instruments with a Zernike phase-mask sensor. II. Concept validation with ZELDA on VLT/SPHERE," *Astron. Astrophys.* **592**, A79 (2016).
10. M. N'Diaye et al., "Calibration of quasi-static aberrations in exoplanet direct-imaging instruments with a Zernike phase-mask sensor," *Astron. Astrophys.* **555**, A94 (2013).
11. R. A. Gonsalves, "Phase retrieval and diversity in adaptive optics," *Opt. Eng.* **21**(5), 215829 (1982).
12. M. Lamb et al., "Non-common path aberration corrections for current and future AO systems," *Proc. SPIE* **9148**, 914857 (2014).
13. R. G. Paxman, T. J. Schulz, and J. R. Fienup, "Joint estimation of object and aberrations by using phase diversity," *J. Opt. Soc. Am. A* **9**, 1072–1085 (1992).
14. M. Lamb et al., "Estimating phase errors from pupil discontinuities from simulated on sky data: examples with VLT and Keck," *Proc. SPIE* **9909**, 99096D (2016).
15. R. Conan and C. Correia, "Object-oriented Matlab adaptive optics toolbox," *Proc. SPIE* **9148**, 91486C (2014).
16. L. M. Mugnier, A. Blanc, and J. Idier, "Phase diversity: a technique for wave-front sensing and for diffraction-limited imaging," *Adv. Imaging Electron. Phys.* **141**, 1–76 (2006).
17. J. Lagarias et al., "Convergence properties of the Nelder–Mead simplex method in low dimensions," *SIAM J. Optim.* **9**(1), 112–147 (1998).
18. B. Paul et al., "Compensation of high-order quasi-static aberrations on SPHERE with the coronagraphic phase diversity COFFEE," *Astron. Astrophys.* **572**, A32 (2014).
19. O. Herscovici-Schiller, L. M. Mugnier, and J. F. Sauvage, "An analytic expression for coronagraphic imaging through turbulence. Application to on-sky coronagraphic phase diversity," *Mon. Not. R. Astron. Soc. Lett.* **467**(1), L105–L109 (2017).
20. A. S. Jurling and J. R. Fienup, "Extended capture range for focus-diverse phase retrieval in segmented aperture systems using geometrical optics," *J. Opt. Soc. Am. A* **31**, 661–666 (2014).
21. L. Jollissaint et al., "Retrieving the telescope and instrument static wavefront aberration with a phase diversity procedure using on-sky adaptive optics corrected images," *Proc. SPIE* **8447**, 844716 (2012).
22. S. Meimon, T. Fusco, and L. M. Mugnier, "Lift: a focal-plane wavefront sensor for real-time low-order sensing on faint sources," *Opt. Lett.* **35**(18), 3036–3038 (2010).
23. C. L. Mehta, "Determination of spectral profiles from correlation measurements," *Nuovo Cimento* **36**, 202–205 (1965).
24. L. M. Mugnier et al., "On-line long-exposure phase diversity: a powerful tool for sensing quasi-static aberrations of extreme adaptive optics imaging systems," *Opt. Express* **16**, 18406–18416 (2008).
25. R. W. Goodrich, "NIRC-2 Observer's Manual," Version 1.1, 2002, <http://www2.keck.hawaii.edu/inst/nirc2/Manual/ObserversManual.html#Section2.2> (2016).

Masen P. Lamb is currently a PhD candidate in the Department of Physics and Astronomy at the University of Victoria, Canada. He received his BSc degree in astronomy at the University of British Columbia, Canada. His current research involves the use of adaptive optics (AO) to resolve metal-poor stars in dense populations such as the Galactic Centre. His research also focuses on the calibration of noncommon path aberrations in contemporary AO systems.

Biographies for the other authors are not available.