# Low-wind-effect impact on Shack-Hartmann-based adaptive optics. Partial control solution in the context of SPHERE and GRAVITY+

N. Pourré, J. -B. Le Bouquin, J. Milli, J.-F. Sauvage, Thierry Fusco, C. Correia, S. Oberti

A&A 665, A158 (2022)
https://doi.org/10.1051/0004-6361/202243432


Astronomy & Astrophysics

# Low-wind-effect impact on Shack–Hartmann-based adaptive optics

## Partial control solution in the context of SPHERE and GRAVITY+

N. Pourré[1], J.-B. Le Bouquin[1], J. Milli[1], J.-F. Sauvage[2,3], T. Fusco[2,3], C. Correia[4], and S. Oberti[5]

[1] Univ. Grenoble Alpes, CNRS, IPAG, 38000 Grenoble, France
e-mail: nicolas.pourre@univ-grenoble-alpes.fr
[2] DOTA, ONERA, Université Paris Saclay (COmUE), France
[3] Aix Marseille Univ, CNRS, CNES, LAM, Laboratoire d'Astrophysique de Marseille, Marseille, France
[4] Space ODT – Optical Deblurring Technologies, Rua Direita de Francos, 1021, Rés-Do-Cháo Esquerdo 4250-194 Porto, Portugal
[5] European Southern Observatory, Karl-Schwarzschild-str-2, 85748 Garching, Germany

Received 28 February 2022 / Accepted 18 June 2022

**ABSTRACT**

*Context.* The low wind effect (LWE) occurs at the aperture of 8-meter class telescopes when the spiders holding the secondary mirror get significantly cooler than the air. The effect creates phase discontinuities in the incoming wavefront at the location of the spiders. Under the LWE, the wavefront residuals after correction of the adaptive optics (AO) are dominated by low-order aberrations, pistons, and tip-tilts, contained in the pupil quadrants separated by the spiders. Those aberrations, called petal modes, degrade the AO performances during the best atmospheric turbulence conditions. Ultimately, the LWE is an obstacle for high-contrast exoplanet observations at a small angular separation from the host star.
*Aims.* We aim to understand why extreme AO with a Shack-Hartmann (SH) wavefront sensor fails to correct for the petal tip and tilt modes, while these modes imprint a measurable signal in the SH slopes. We explore if the petal tip and tilt content of the LWE can be controlled and mitigated without an additional wavefront sensor.
*Methods.* We simulated the sensitivity of a single subaperture of a SH wavefront sensor in the presence of a phase discontinuity across this subaperture. We explored the effect of the most important parameters: the amplitude of the discontinuity, the spider thickness, and the field of view. We then performed end-to-end simulations to reproduce and explain the behavior of extreme AO systems based on a SH in the presence of the LWE. We then evaluated the efficiency of a new mitigation strategy by running simulations, including atmosphere and realistic LWE phase perturbations.
*Results.* For realistic parameters (i.e. a spider thickness at 25% of a SH subaperture, and a field of view of $3.5\lambda/d$), we find that the sensitivity of the SH to a phase discontinuity is dramatically reduced, or even reversed. Under the LWE, a nonzero curl path is created in the measured slopes, which transforms into vortex-structures in the residuals when the loop is closed. While these vortexes are easily seen in the residual wavefront and slopes, they cannot be controlled by the system. We used this understanding to propose a strategy for controlling the petal tip and tilt modes of the LWE by using the measurements from the SH, but excluding the faulty subapertures.
*Conclusions.* The proposed mitigation strategy may be of use in all extreme AO systems based on SH for which the LWE is an issue, such as SPHERE and GRAVITY+.

**Key words.** instrumentation: high angular resolution – instrumentation: adaptive optics

## 1. Introduction

The Spectro-Polarimetric High-Contrast Exoplanet REsearch (SPHERE) instrument at the Very Large Telescope (VLT) is a high-contrast imager optimized for exoplanet hunting (Beuzit et al. 2019). The instrument is equipped with extreme adaptive optics (AO) system (called SAXO, Fusco et al. 2016; Sauvage et al. 2016b) that routinely reaches a Strehl ratio (SR) of 90% in the $H$ band. SPHERE has three scientific arms with different detectors: ZIMPOL (Schmid et al. 2018), allowing for polarimetric observations in the optical, IRDIS (Dohlen et al. 2008), for dual-band imaging and spectroscopy in the near-infrared, and IFS (Claudi et al. 2008), an integral field spectrograph in the near-infrared. Also, the instrument includes cutting-edge coronagraphs at both optical and near-infrared wavelengths, which enable a final contrast of $10^{-5}$ in the $H$ band for exoplanet observations at a 500 mas separation from the host star (Langlois et al. 2021; Mouillet et al. 2018). Since 2014, SPHERE has achieved groundbreaking observations in the field of exoplanets (e.g., Chauvin et al. 2017; Keppler et al. 2018; Vigan et al. 2021) and circumstellar disks (e.g., van Boekel et al. 2017; Milli et al. 2017; Boccaletti et al. 2020; Ginski et al. 2021).

However, under the best atmospheric conditions, the instrument performances are hampered by the so-called low wind effect (LWE). This effect has been observed since the commissioning of SPHERE in 2014 and it was highlighted very early on as a major limitation of the instrument (Sauvage et al. 2015). The LWE is responsible for differential aberrations (petal modes) between the quadrants of the unit telescope (UT) pupil separated by the four spiders that hold the secondary mirror. Measurements

with the Zernike phase mask ZELDA (N'Diaye et al. 2013, 2016) under LWE conditions on SPHERE have shown typical petal-pistons (hereafter PPs) and petal-tip-tilts (PTTs) in the residual phase screens after the AO correction. The amplitude of the uncorrected aberrations measured with ZELDA ranges from 600 to 800 nm peak-to-valley (ptv) optical path difference (OPD; Sauvage et al. 2015). As a result, the focal plane images are affected by bright side lobes at the location of the first Airy ring responsible for Strehl losses, typically around 30% in the H band (Milli et al. 2018). On a coronagraphic image, the LWE ruins the contrast by a factor of up to 50 at a separation of 0.1″. Unfortunately, the LWE is not restricted to SPHERE, but is also observed with the Adaptive Optics Facility (AOF, Oberti et al. 2018) at the VLT and with SCExAO at Subaru (Bos et al. 2020).

The commonly acknowledged physical explanation for the LWE is the following. At night, the spiders holding the secondary mirror radiate their heat to the clear sky and their temperature drops to 2 ∼ 3 °C below the ambient air temperature. When the wind at the top of the telescope dome drops below 5 m s$^{-1}$, the air in the dome is not well mixed and a laminar flow can develop around the spiders. Due to thermal exchange, the air efficiently cools down near the spiders, generating layers of colder air in the vicinity of the spiders (Holzlöhner et al. 2021). Ultimately, temperature differences are responsible for optical index differences on each side of the spider, and therefore the discontinuity of the OPD.

A passive mitigation was applied in 2017 on UT3 where SPHERE operates. It consisted in applying a coating on the spiders to limit their thermal emissivity in the mid-infrared. The occurrence frequency of the LWE on SPHERE dropped from ∼20% of the observing time to ∼3.5% of the time (Milli et al. 2018). Still, the LWE continues to degrade the observations when atmospheric conditions are the best. This explains the recent efforts to develop additional mitigation strategies to control the LWE (see Vievard et al. 2019, for an overview of focal plane wavefront sensor strategies). The focal plane wavefront sensing solution called Fast & Furious (F&F, Keller et al. 2012; Korkiakoski et al. 2012, 2014) is one of the most advanced algorithms for LWE control and it has been successfully tested in the laboratory (Wilby et al. 2018) and on-sky using Subaru/SCExAO (Bos et al. 2020). It uses sequential phase diversity (Gonsalves 2002) from an additional focal-plane sensor to measure the PP and PTT aberrations, and the first 50 Zernike modes for non-common path aberrations control. Here, we propose following a complementary approach: understanding the observed behavior of the AO under the LWE in order to propose improvements without requiring a new sensor.

The first purpose of this paper is to understand why the extreme AO of SPHERE fails to correct the PTT aberrations induced by the LWE. Indeed, the PTTs are normal tip-tilts (TTs) on most parts of the pupil. As such, they imprint a recognizable pattern in the slopes measured by the Shack-Hartmann (SH) wavefront sensor. The second purpose is to propose and evaluate a mitigation strategy using, as much as possible, the information already provided by the SH. The paper uses low-level and end-to-end simulations to achieve these goals, and is organized as follows. In Sect. 2, we investigate how bad SH measurements for phase discontinuities lead to uncorrected aberrations. We start in Sect. 2.1 with a low-level study of the slope measured by a single SH subaperture exposed to a phase discontinuity instead of a smooth phase slope. Then, in Sect. 2.2, we link these results with the residuals of an end-to-end AO simulation, and finally reproduce post-AO residuals observed on SPHERE. Section 3 is dedicated to the description of a mitigation algorithm that makes use of the SH information to reliably measure and control the PTT modes. The paper concludes with discussions in Sect. 4 about the interest and limitations of the proposed mitigation, considering the known properties of the LWE.

In the paper, nonbold variables are scalars, bold variables are vectors, and bold-underlined variables are matrices. The symbol $*$ is for element-wise multiplication and $\cdot$ is for matrix product.

## 2. From bad wavefront measurements to uncorrected aberrations

The LWE induces aberrations that are not corrected by the adaptive optics. In this section we describe how bad measurements by the SH can lead to strong low-order post-AO residuals.

### 2.1. Shack-Hartmann's sensitivity to a phase discontinuity

The most problematic features of the residual aberrations created by the LWE are the sharp phase discontinuities along the spiders. First, we investigate how such a phase discontinuity affects the measurement from a single SH subaperture. Similar studies have already been performed in the context of detecting phasing errors between primary mirror segments of the W.M. Keck Observatory (Chanan et al. 1998, 2000; van Dam et al. 2017). The conclusion is that, in the weak phase regime, a SH should be able to measure a phase jump. Here, we go further by investigating the influence of the following parameters: the position and amplitude of the discontinuity, the field of view, and the presence of a spider.

We built a basic simulation to propagate the electric field from the pupil plane to the focal plane of the subaperture:

$$\mathbf{I} = |\,\mathscr{F}\{\mathbf{A}\exp(i\boldsymbol{\phi})\}\,|^2, \tag{1}$$

where $\mathbf{I}$ is the intensity at the focus of the subaperture, $\mathscr{F}$ is the Fourier transform operation, $\mathbf{A}$ is the pupil transmission (e.g., a simple square of size $d$ for an unobstructed subaperture), and $\boldsymbol{\phi}$ is the phase screen in front of the subaperture. We used a sampling of ≈200 point across the subaperture plane and ≈10 points per $\lambda/d$ in the image plane. Tests with a finer sampling resulted in no significant changes in the results. We used a center of gravity (CoG) calculation to determine the spot position in the $x$ direction:

$$\mathrm{CoG} = \frac{\int x\,\mathbf{I}\,\mathrm{d}x}{\int \mathbf{I}\,\mathrm{d}x}. \tag{2}$$

The integration for the center of gravity was limited to a given field of view, which was a free parameter.

The expected displacement of the spot in units of $\lambda/d$ due to a phase slope of ptv amplitude $\Delta\phi = a \times 2\pi$ across the subaperture is given by CoG = $a$. This relationship does not always hold true, especially when the phase is not a gentle slope, but instead contains a discontinuity. It is the purpose of the following simulations to quantify the extent to which the CoG measurement deviates from the expected value, and under which conditions.

#### 2.1.1. Effect of position and field of view

We first analyze how the measurement is affected by the position of the discontinuity and the field of view used to compute the CoG. We restrict this analysis to the weak-phase regime. Results are given in Fig. 1. When the field of view is wide (e.g., 40 $\lambda/d$), the CoG properly estimates the discontinuity, unless when the

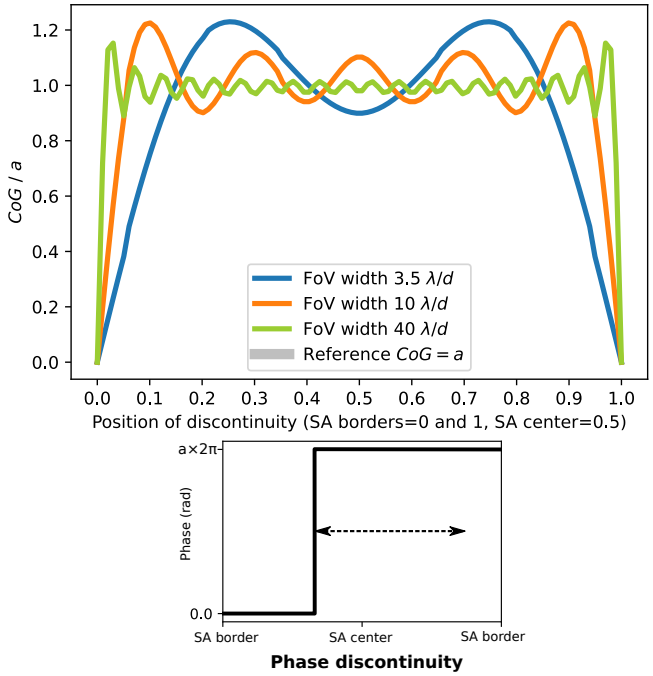


**Fig. 1.** Impact of the position of the phase discontinuity on the CoG measurement. *Top:* sensitivity to a phase discontinuity with respect to the position of the discontinuity in the subaperture (SA) for different fields of view (FoV). The simulation is carried out for $a = 0.05$ (weak-phase regime). *Bottom:* 1D illustration of the free parameter, the position of the discontinuity.

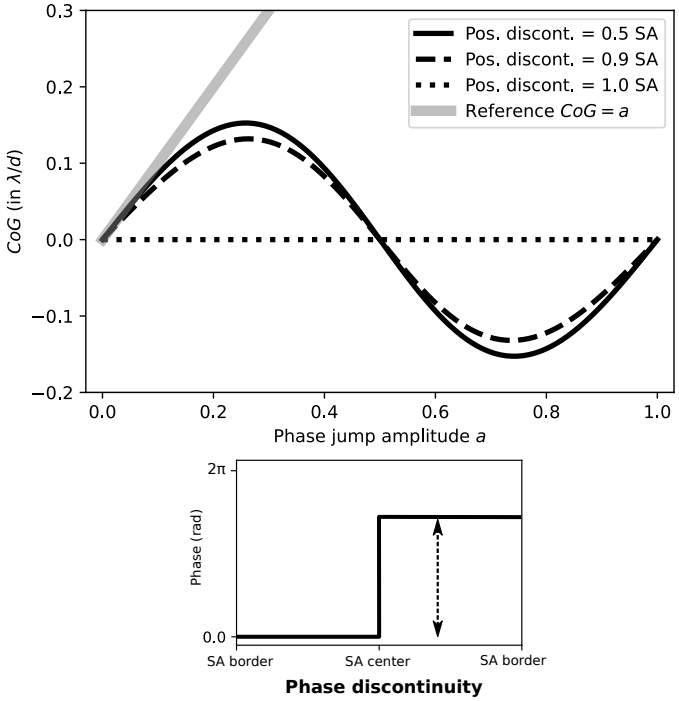


**Fig. 2.** Impact of the amplitude of the discontinuity on the CoG measurement. *Top*: sensitivity to a phase discontinuity with respect to the amplitude of the discontinuity, where a = 1 corresponds to a $2\pi$ rad phase discontinuity. The field of view is restricted to a 3.5 $\lambda/d$ width. Black lines correspond to different positions of the discontinuity in the pupil plane (0.5 is for the center, 0 and 1 are for the edges). The gray line corresponds to the reference value CoG = $a$. *Bottom*: 1D illustration of the free parameter, the amplitude of the discontinuity.

discontinuity is located very close to the edge of the subaperture. This is an expected result, as already pointed out in van Dam et al. (2017). However, for a field of view realistically restricted to 3.5 $\lambda/d$, the sensitivity significantly depends on the position of the discontinuity. The CoG estimate can be erroneous by up to 20% with respect to the expected value.

### 2.1.2. Effect of the amplitude

We then analyze the response when the discontinuity amplitude goes beyond the weak-phase regime. The results are presented in Fig. 2. As expected, the CoG measurement evolves nonlinearly with respect to the amplitude of the discontinuity, and wraps with a period of $2\pi$. To put it simply, for an amplitude larger than $\pi/2$ rad ($a \geq 0.25$), there is no hope for the CoG to properly estimate the amplitude of the discontinuity. The relationship with the position of the discontinuity is consistent with the sensibility curves in Fig. 1 for a field of view of 3.5 $\lambda/d$. In fact, we found that the effect of amplitude is decoupled from the effect of field of view and position. That is, all configurations can be estimated quantitatively by scaling the results of Fig. 1 with the results of Fig. 2.

### 2.1.3. Effect of the spider thickness

In practise, phase discontinuities occur at the location of the spiders. It is therefore necessary to investigate how the CoG is affected by a partial obscuration in the subaperture. The results are shown in Fig. 3. The sensitivity to the discontinuity decays dramatically with the thickness of the spider. Moreover, increasing the field of view does not provide a remedy for the missing sensitivity. For a spider size of 25% of the subaperture width, and a field of view of 3.5 $\lambda/d$, the CoG provides a measurement with a sign opposite to the applied perturbation.

One could question whether the issue of partial illumination also affects the measurement of slopes for a continuous wavefront, for example when sensing the turbulence without any LWE. Figure 3 shows that the loss in sensitivity is significantly less when considering a phase slope instead of a phase discontinuity. For a field of view of 3.5 $\lambda/d$, the loss of sensitivity reaches around 50%. Still, the CoG always provides an estimate with the correct sign. For wider fields of view, the loss in sensitivity due to the spider for the phase slope becomes less prominent.

### 2.1.4. Application to SPHERE

The SPHERE instrument of the VLT uses a field of view of 3.5 $\lambda/d$ to compute the CoG. The 5cm thick spiders of the VLT block 25% of a subaperture (40 subapertures across the 8 m pupil). Unfortunately, Fig. 3 shows that it is the worst configuration, that is to say, it is the configuration for which the CoG measurement of a discontinuity is the most different from the expected value. Figure 4 summarizes the resulting effect on the sensitivity around the spider. For a phase discontinuity, the CoG estimate never gets close to the expected value. At best, the sensitivity is 0.1. At worst, the sensitivity reaches –0.5 when the discontinuity is at the center of the subaperture. This result provides an explanation for the "contra-moving spots" observed on SPHERE (SPHERE commissioning report, Sauvage, private communication). Obviously, one expects such reversed sensitivity to have a dramatic effect on the closed-loop operation. For a phase slope, the sensitivity losses are less severe

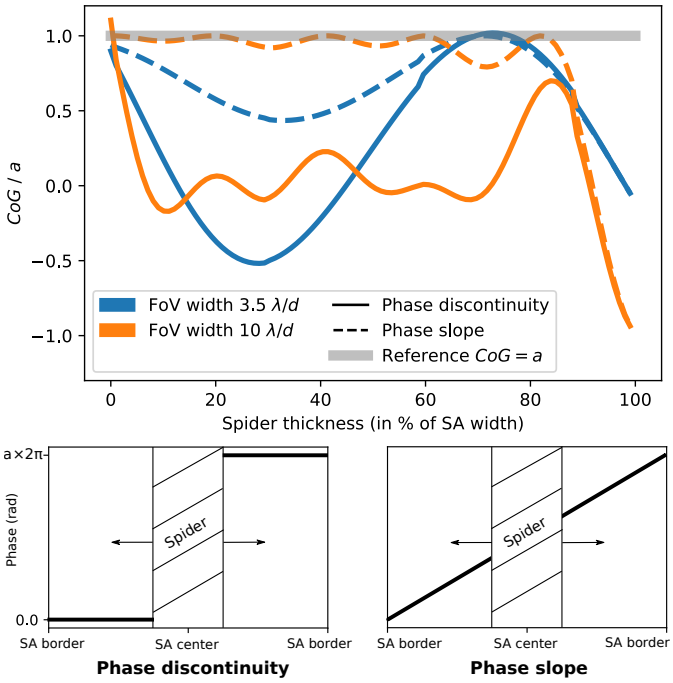


**Fig. 3.** Impact of the spider thickness on the CoG measurement. *Top:* sensitivity to a phase discontinuity (solid lines) and to a phase slope (dashed lines) with respect to the spider thickness. The simulation is carried out with $a = 0.05$ (weak phase regime). The spider obstructs the subaperture (SA) at the center of the subaperture. The spider thickness varies from 0% (infinitely thin) to 100% of the SA width (full obstruction). Two different fields of view (FoV) are tested. *Bottom:* 1D illustration of the free parameter, the spider thickness, in both the phase discontinuity (*left*) and the phase slope (*right*) cases.

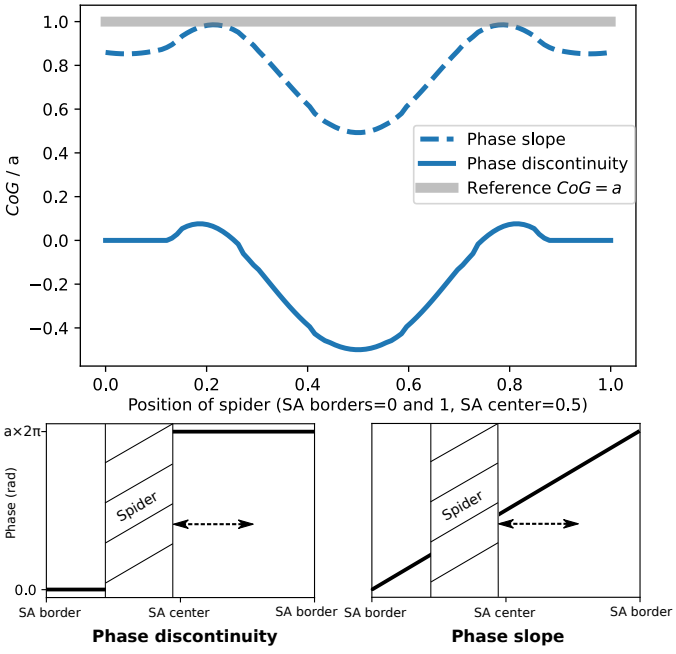


**Fig. 4.** Impact of the spider position on the CoG measurement. *Top:* sensitivity to a phase discontinuity and a phase slope with respect to the position of the spider in the subaperture (SA). The spider thickness is 25% of the SA width and the field of view is 3.5 $\lambda/d$. The simulation is carried out for $a = 0.05$ (weak phase regime). *Bottom:* 1D illustration of the free parameter, the position of the spider, in both the phase discontinuity (*left*) and the phase slope (*right*) cases.

than for the discontinuity; at worst the sensitivity drops to 0.5, but always keeps the correct sign. As long as this sensitivity error remains within the gain margin of the controlled modes, it will be ultimately corrected by the feedback loop.

Overall, the different behavior of the CoG when exposed to a phase discontinuity or a phase slope explains why the presence of spiders in the aperture is not problematic for measuring continuous atmospheric aberrations, but is an issue for measuring discontinuous LWE aberrations.

### 2.2. Uncorrected aberrations

We expect the bad SH measurement to have an impact on the AO correction. In this section we use end-to-end AO simulations to characterize the uncorrected aberrations and understand why they arise.

#### 2.2.1. Design of the adaptive optics simulation

We used the HCIPy (High-Contrast Imaging for Python, Por et al. 2018) AO simulator to model a high-order AO system. This tool enables a simulation for the spots of each subaperture in the presence of phase discontinuity thanks to a proper treatment of the optical propagation. Electric fields are sampled with 480x480 grid points over the full pupil. Finer sampling resulted in no significant change.

The overall design is similar to SPHERE. The SH wavefront sensor operates at $\lambda = 700$ nm and has $40 \times 40$ subapertures in a Fried configuration with respect to the deformable mirror (DM). The DM is composed of 1377 actuators (41 actuators per diameter). The system controls the first 990 Karhunen–Loève (KL) modes (piston excluded). These KL modes are defined by a **K2DM** (KL to DM) $1377 \times 990$ matrix. We calibrate the interaction matrix **K2S** (KL to slopes) and take the pseudo-inverse to obtain the reconstruction matrix **S2K**. All modes are controlled with the same leaky integrator (leak $l = 0.01$, and gain $g = 0.3$):

$$\mathbf{DM}_{t+1} = \underline{\mathbf{K2DM}} \cdot [\,(1 - l)\,\mathbf{K}_t + g\,\underline{\mathbf{S2K}} \cdot \mathbf{S}_t\,]. \quad (3)$$

The circular telescope pupil has a diameter of 8 m and a central obscuration whose diameter is 1.116 m. The pupil is segmented into four quadrants. We apply a flux criterion to discard the subapertures outside the useful pupil, setting the threshold at 50% of the flux received by a nonobstructed aperture. This selection criterion keeps a total of 1160 subapertures and always keeps the subapertures located behind the spiders.

#### 2.2.2. Simplified perfectly blind configuration

We set up a simplified, symmetric pupil, where the junction between quadrants were aligned with the SH grid and passed between neighboring subapertures. For this configuration, the AO response to a simple PP and a simple PTT perturbation is shown in Fig. 5.

Intuitively, and as seen in Sect. 2.1, this configuration is perfectly blind to discontinuities between quadrants that are pure PPs (top). Yet, the $40 \times 40$ SH measures very small nonzero slopes around discontinuities. This comes from the leaking of the diffraction pattern of each subaperture in the field of view of the neighboring subapertures. The subapertures on each side of the discontinuity differ by a piston. Therefore, the diffracted electric fields interfere and slightly bias the CoG positions. This effect explains the small commands applied to the DM and the subsequent small reduction of the wavefront residuals (0.93). It should

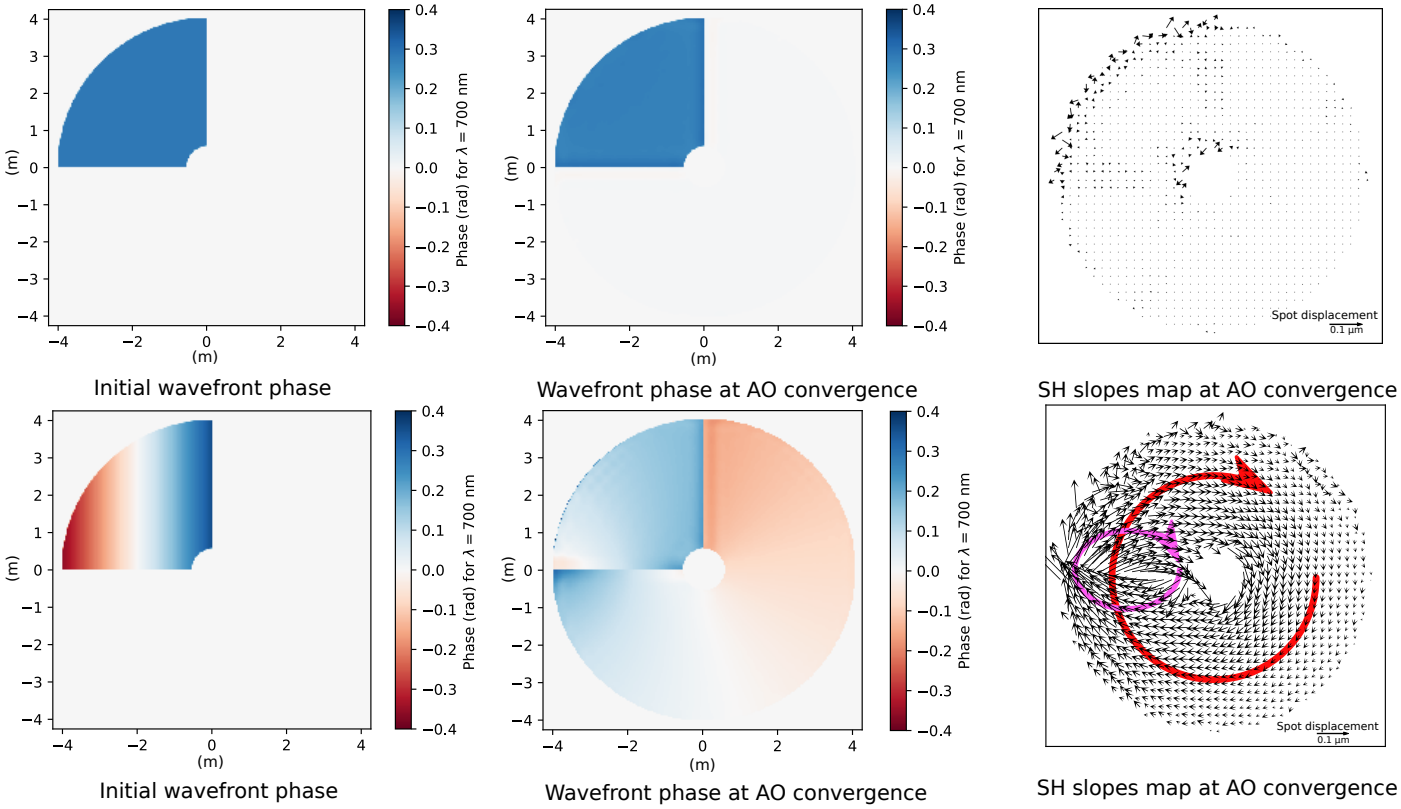


**Fig. 5.** End-to-end simulation of an AO loop in the presence of discontinuities. Discontinuities pass exactly between SH subapertures. The spider is infinitely thin, and there is no atmosphere. *Top:* static PP of a 32 nm ptv (weak-phase regime) on one quadrant. The ratio between the rms at convergence and the initial rms is 0.93. There are almost no signals in the residual SH slope map. *Bottom:* static PTT of an 80 nm ptv (weak-phase regime) on one quadrant. The ratio between the rms at convergence and the initial rms is 1.02. The red vortex in the residual slopes comes from the nonzero curl of the initial wavefront over a path circling around the central obscuration. The pink vortex in the residual slopes comes from the nonzero curl of the initial wavefront over a smaller path centered in $(-2, 0)$.

be noted that the slope map at the AO convergence contains no global residuals; all that is seen has been corrected.

The situation is different for the PTT perturbation (bottom). The aberration is tentatively corrected, but the applied correction leaks into the entire pupil. The residual wavefront is a vortex phase that shows two prominent curl patterns in the map of residual slopes measured by the SH. Two questions arise from this result, and we thus sought to understand what creates the curl patterns, and why these patterns were not corrected.

The answer to the first question is the loss in sensitivity. This is illustrated in the top panel of Fig. 6. The SH provides an erroneous estimate of the discontinuity amplitude. Therefore, the integration of the slopes along a closed path crossing the discontinuity only once is nonzero. There is an interesting explanation for the double vortex observed in the residual slopes of the bottom panel of Fig. 5. The red vortex in the residual slopes comes from the nonzero curl of the initial wavefront over a path circling around the central obscuration, which is going through $(-2, 0)$, $(0, +2)$, $(+2, 0)$, and $(0, -2)$. This path crosses a discontinuity only once, at $(0, +2)$. The pink vortex in the residual slopes comes from the nonzero curl of the initial wavefront over a smaller path going through $(-3, 0)$, $(-2, +1)$, $(-1, 0)$, and $(-2, -1)$. This path crosses two discontinuities, once at $(-3, 0)$ and once at $(-1, 0)$, but accumulating the losses in the same direction.

The answer to the second question is that curl patterns are out of the control space. It is well known that the curl of the gradient of a scalar differentiable field is zero. The DM being a continuous, smooth surface, it can only create modes without curl. Even the best approximation of a vortex phase map that the DM can create still has a zero curl. This is illustrated by the bottom panel of Fig. 6. This property also applies to the KL modes that form a subspace of the DM space. It explains why any nonzero curl patterns in the residual slopes are not corrected by the AO.

### 2.2.3. Dependence with the position of the discontinuity

From Sect. 2.1, we expect that the convergence state of the AO loop depends on the exact position of the discontinuity and the possible thickness of the spider. We verified these behaviors by running end-to-end simulations with varying these parameters. The results are presented in Appendix A. The performance of the AO loop is quantified as the ratio between the rms of the residual wavefront at convergence and the rms of the input aberration. The results match the predictions from Fig. 1 (the case with the infinitely thin spider) and Fig. 4 (the case with the 25% thick spider). The correction is best for the setups where the sensitivity is close to one, and worst when the sensitivity has the wrong sign.

This validates our understanding of the link between the local effect (reduced, reversed sensitivity of the discontinuity) and the global effect (unseen piston modes and uncorrectable curl modes), and allows us to make some quantitative conclusions. We find that a SH with thin spiders crossing the subapertures at an adequate position (0.4) could deal with discontinuities up

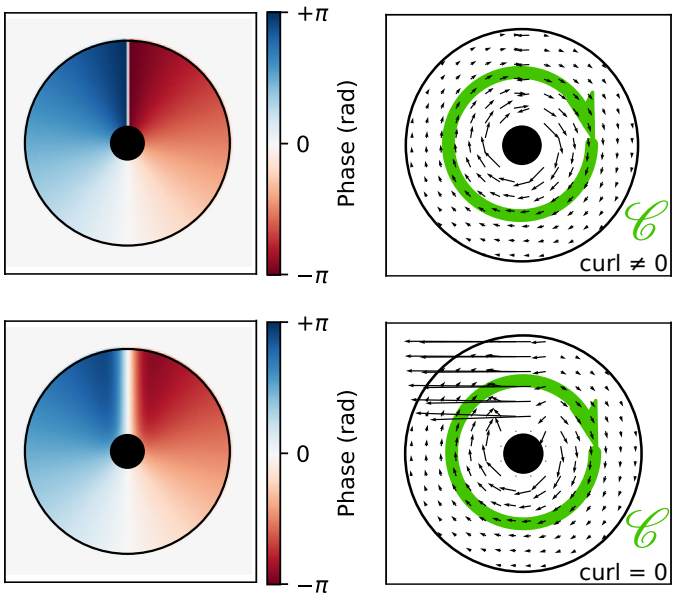


**Fig. 6.** Schematic for the gradient measured by a SH for a sharp discontinuity (*top*) and for a smoother pattern (*bottom*). In the top case, the amplitude of the discontinuity is not estimated correctly, and consequently the sum of SH slopes along a closed path $\mathscr{C}$ is not equal to zero. In the bottom case, the sum of slopes along the path $\mathscr{C}$ is equal to zero because the strong slope in affected subapertures compensate exactly for the rest of the path.

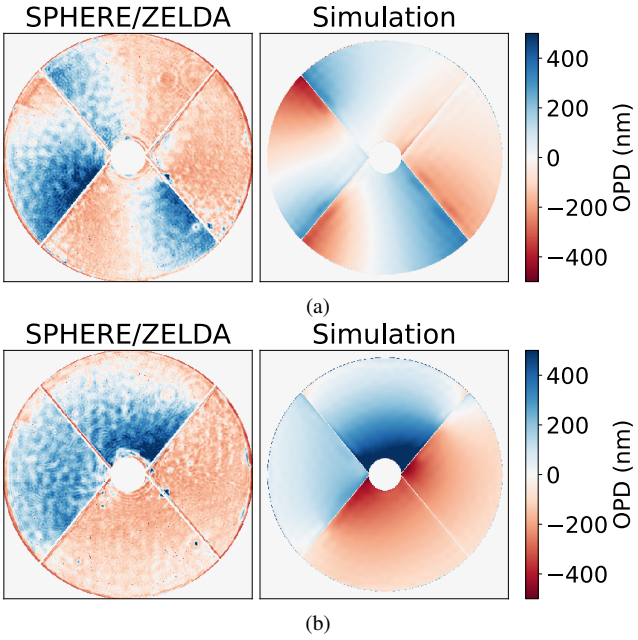


**Fig. 7.** Qualitative comparison between ZELDA measurements (*left*) and simulated post-AO residuals (*right*) for two LWE events (*a*) and (*b*). The ZELDA measurements have a 1 second detector integration time. They were taken during the night of 2014 October 8. The simulations match the typical SPHERE configuration. The result after convergence of the loop is shown here.

to $\lambda/4$. However, in the presence of a realistic spider obscuration of 25% of a subaperture, the impact of discontinuities is considerable.

#### 2.2.4. Reproduction of SPHERE low-wind-effect residuals

We modified the end-to-end simulations to better match the configuration of SPHERE. We used a realistic VLT pupil with 5 cm thick spiders and with the proper angle for each spider. In this realistic configuration, the spiders (and the discontinuities) cross the SH subapertures at various positions and with different angles. We implemented a spatial filter of $2\lambda/d$ in front of the SH ($d$ being the SH subapertures' width) and a Gaussian weighting on the SH focal-plane spots. We offloaded the TT control on a dedicated TT mirror. We implemented the differential tip-tilt sensor (Baudoz et al. 2010), whose aim is to maintain the centering of the star, behind the coronagraph by measuring, at a slower frequency, the actual position of the star at the focal plane in the H band. In all following simulations, the differential tip-tilt sensor frequency was 1 Hz and the AO loop frequency was 1.2 kHz.

We found that it is interesting to include the differential tip-tilt sensor in our simulation. Indeed, this sensor measurement is a CoG in the H-band focal-plane image. It is sensitive to petal perturbations (e.g., PPs and PTTs) that project on the global TT but are unseen by the SH. Yet, we have shown in Sect. 2.1 that phase discontinuities across an aperture can bias CoG measurements. This conclusion is also true for the differential tip-tilt sensor that makes an erroneous estimate in the case of a discontinuous wavefront, thus enforcing a global TT in the residual phase maps. Simulating the differential tip-tilt sensor is necessary in order to achieve residuals with the same overall structures as the ones observed with the ZELDA sensor, especially the residual global TT.

We used this realistic simulation of SPHERE AO. The LWE was simulated by an injected perturbation in the pupil plane combining discontinuous PP perturbations and discontinuous PTT perturbations. This way, we were able to reproduce some typical PTT and PP structures observed in the post-AO residuals of the instrument (Fig. 7). Figure 7a displays a residual pattern dominated by global vortexes (around the central obscuration) and local vortexes (around the center of individual spiders). Figure 7b displays a residual pattern dominated by strong PPs.

## 3. Shack–Hartmann-assisted low-wind-effect control

In a weak scintillation regime, the aberration created by the atmospheric turbulence is a continuous and differentiable scalar field (no branch points or branch cuts, Primmerman et al. 1995; Fried 1998). Optical vortexes are thus generally neglected in space-to-ground AO. The problem of vortex reconstruction and branch point identification in the slopes measured by a SH has, however, driven important literature in other contexts (Fried 1998; Tyler 2000; Luo et al. 2015; Wu et al. 2021). Here, we propose building on the successful approaches that made use of petaling modes to correct for the LWE (Sauvage et al. 2016a; Wilby et al. 2018; N'Diaye et al. 2018). These works rely on an additional wavefront sensor to control the petaling modes, generally focal-plane images. Here, we wish to utilize the information encoded inside the nonzero residual slope pattern of the SH sensor as much as possible.

### 3.1. Description of the algorithm to mitigate the low wind effect

We designed an algorithm that measures the PTTs from the SH and the PPs from a *H* band focal plane. The corrective command

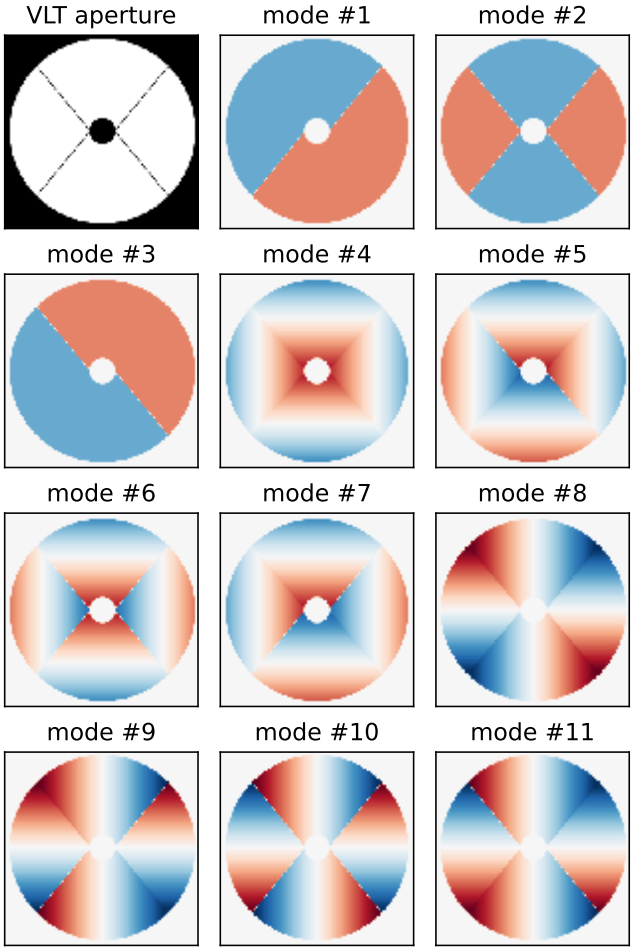


**Fig. 8.** 11 orthogonal LWE modes of the VLT pupil used in this study as pupil-plane perturbations. This compares to the classical decomposition shown in Fig. 6 of Sauvage et al. (2016a). Modes #1 to #3 are PP modes. Modes #4 to #11 are PTT modes. Odd numbers are for odd modes and even numbers are for even modes. Modes #4 and #8 do not contain discontinuities, and are removed from the control basis. Mode #10 controls the global curl pattern seen in Fig. 5.

is given to the deformable mirror via a modification of the reference slopes.

#### 3.1.1. An alternate modal basis for the low wind effect

The classic basis used to describe the LWE aberrations is composed of 11 modes: the PP and PTT modes in each of the four quadrants of the pupil (three PP modes and eight PTT modes, Sauvage et al. 2016a). We propose decomposing this basis into a new set of odd and even modes (Fig. 8). There are two advantages to this. First, expressing aberrations as odd and even modes is more convenient for a focal-plane analysis, if needed. Second, two of the new PTT modes do not contain phase discontinuities (modes #4 and #8 in Fig. 8). Even if they are not differentiable at the junction between quadrants, those two modes are properly handled by the AO loop. This new basis reduces the number of PTT modes to control from eight to six. To be explicit: modes #4 and #8 are not included in the control matrices, but they are included in the description of the input LWE perturbation.

In practise, the modes of Fig. 8 are not perfectly realizable by the DM. We defined the matrices $\underline{\mathbf{PP2K}}$ and $\underline{\mathbf{PTT2K}}$, which are the best approximations of the PP and PTT modes, in phase space, as per the KL basis. We arbitrarily restricted the approximation to the KL space and not the full DM space in order to remain inside the control space of the high-order AO loop. It is important to note, however, that we have not demonstrated that this is a strict requirement. The synthetic interaction matrix of the three PP modes and of the six remaining PTT modes can be constructed from the already known interaction matrix of the KL modes:

$$\underline{\mathbf{PTT2S}} = \underline{\mathbf{K2S}} \cdot \underline{\mathbf{PTT2K}}, \tag{4}$$

$$\underline{\mathbf{PP2S}} = \underline{\mathbf{K2S}} \cdot \underline{\mathbf{PP2K}}. \tag{5}$$

A naive approach is to simply add these modes into the control modal basis. This has been tried in the existing instrument (Sauvage, priv. comm.), and we reproduced the experiment in our simulation. This is doomed to fail, because the issue is not the completeness of the control space. Instead, as demonstrated in previous sections, the issue arises from improperly seen pistons and curl modes. For the PPs, there is nothing we can do with the SH. For the curl modes, it is possible to modify the measurement space in order to improve their visibility by the system.

#### 3.1.2. Measuring the petal-tip-tilt modes with a Shack-Hartmann

The PTT modes are wrongly corrected by the system because the interaction matrix of their best approximation (without curl) does not match their actual imprint in the signal (with curl). This is the effect discussed in previous sections and illustrated in Fig. 6. One way to remedy this is to restrict the measurement space to the subspace for which there is a correct match between the expectation and the actual signal. This subspace is simply made of the subapertures that are not affected by the discontinuities.

On the one hand, discarding too many continuous subapertures will lead to the so-called island effect. It is an issue on the Extremely Large Telescope where spiders are too thick to ensure a continuity between neighboring sections of the pupil (Schwartz et al. 2017; Hutterer et al. 2018; Bertrou-Cantou et al. 2020). We reproduced this unwanted behavior in our simulations when discarding the subapertures partially blocked by the spiders. On the other hand, we have shown in Sect. 2.1.4 that the turbulence is properly seen by subapertures partially blocked by the spiders. Therefore, we decide to use a combination of the two approaches: closing the fast AO loop with all subapertures and all KL modes, to efficiently fight the turbulence and avoid its coupling into the island effect; and controlling specifically the PTT at a slower temporal bandwidth, using only the subaperture not affected by the spiders.

We computed the corresponding reconstruction matrix by restricting the pseudo-inverse $\underline{\tilde{\mathbf{S}}\mathbf{2PTT}} = (\underline{\mathbf{PTT2}\tilde{\mathbf{S}}})^{-1}$ to the subapertures that were not crossed, even minimally, by the spiders (see Fig. 9). The tilde in the above equation indicates that the matrix was restricted to selected subapertures. In fact, discarding these subapertures was a way to force curly slope patterns to project onto the controlled modes. Theoretically, it is possible to control the LWE with a higher number of modes than just the PTT (that is, with modes showing curvature across each quadrant). Here, we followed the standard approach of controlling only a small number of flat modes, as expressed in the modified basis of Fig. 8.

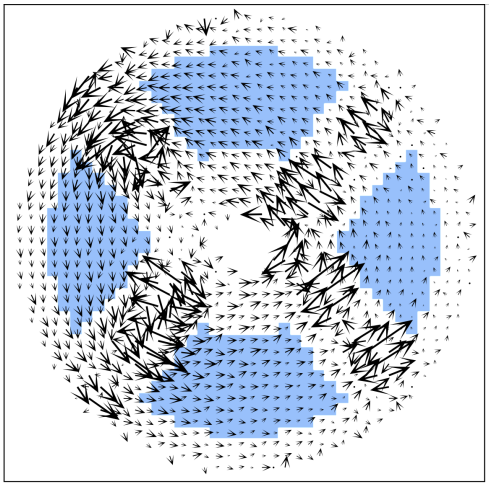

**Fig. 9.** SH slope map. The blue region highlights the subapertures that are used to estimate the PTT modes. Subapertures located close to the spiders, the secondary mirror, and the outer pupil ring are discarded.

#### 3.1.3. Measuring the petal-piston modes

The PP modes only affect the subapertures that are located at the discontinuities, and for which the actual sensitivity is poor and hardly predictable. It is indeed well known that the SH is not a satisfactory sensor for PPs. The only solution is to rely on an additional sensor.

In the simulation, we modified the SPHERE differential tip-tilt sensor in order to measure the three PPs instead of a single, global TT. For this, we used a simple focal-plane analysis inspired by the literature (Korkiakoski et al. 2014; Wilby et al. 2016; Bos et al. 2020). It should be noted that two out of the three PP modes are odd. They could be linearly estimated from a single image measurement, assuming the corresponding interaction matrix had been calibrated. Therefore, out of the 11 initial LWE modes, only one remains to be estimated by a non-linear analysis (mode #2 in Fig. 8). The derivations are detailed in Appendix B. It is not the purpose of this paper to expand on the well-known focal-plane analysis. We implemented it in the simulation to ensure there was no damaging interplay with the proposed solution to control the PTT modes.

#### 3.1.4. Feedback to the adaptive optics loop

By design, the corrections remain within the controlled space of the AO, and would be then flushed out by the closed loop if applied directly on the DM command. The solution is thus to implement the correction via a modification of reference slopes, as was proposed in the very first studies on the LWE on SPHERE (Sauvage et al. 2015, 2016a).

More precisely, the reference slopes $\mathbf{S}^{\mathrm{ref}}$ are modified thanks to simple integrators:

$$\begin{aligned}\mathbf{S}^{\mathrm{ref}}_{t+1} &= \mathbf{S}^{\mathrm{ref}}_{t} - g_{\mathrm{PTT}}\,\underline{\mathbf{PTT2S}}\cdot\underline{\tilde{\mathbf{S}}\mathbf{2PTT}}\cdot\mathbf{S}_t \\ &\quad - g_{\mathrm{PP}}\,\underline{\mathbf{PP2S}}\cdot\mathrm{F}(\mathbf{I}_t),\end{aligned} \tag{6}$$

where F(**I**) is the operation to extract the PP modes out of the focal plan image **I**. For two out of the three PP modes, this operation is a simple matrix multiplication. The scalar coefficients $g_{\mathrm{PTT}}$ and $g_{\mathrm{PP}}$ are the gain for the integration of the PTT and PP commands, respectively. We set $g_{\mathrm{PTT}} = 0.005$ (at the same frame rate than the main loop, 1.2 kHz) and $g_{\mathrm{PP}} = 0.1$ (at a frame rate of 1 Hz). The gain values were empirically chosen to optimize the trade-off between control bandwidth and loop stability.

Figure 10 gives a schematic overview of the complete loop, when incorporating the proposed control for the LWE. We recall that, while the correction is implemented in the slope space, the control space of the LWE is in fact restricted to the KL modes. For the sake of completeness, we also ran the algorithm with the PP and PTT modes expressed on a DM zonal basis instead of the KL basis, and we obtained very similar results.

### 3.2. Simulation parameters

We used the realistic SPHERE simulation described in Sect. 2.2.4. We implemented the proposed control of the PTT modes, as described in 3.1.2. We also modified the differential tip-tilt sensor in order to measure PP modes, as described in Sect. 3.1.3. We included atmosphere turbulence phase screens to verify their possible interplay with the proposed LWE mitigation algorithm. The LWE is only observed under good atmospheric conditions, and we thus simulated such a situation (von Karman power spectral with Fried parameter $r_0 = 16.8$ cm, outer scale $L_0 = 40$ m, coherence time $\tau_0 = 5$ ms, all defined at $\lambda = 500$ nm). We kept the same sequence of atmospheric phase screens in all simulations to permit a direct comparison of the outcome. No noise sources were simulated (no photon noise or read-out noise).

We tested two static LWE phase patterns. The LWE #1 was the input LWE that corresponds to the measured ZELDA post-AO residuals displayed in Fig. 7a. The corresponding input LWE contained a mix of PP and PTT perturbations that displayed a clear curl structure. The simulated post-AO residuals, without any mitigation of the LWE, had 173 nm rms and 650 nm ptv OPD. The LWE #2 was an input LWE that induced very strong PTT in post-AO residuals. The simulated post-AO residuals, without any mitigation of the LWE, had 276 nm rms and 1350 nm ptv OPD. The post-AO residuals for these two LWE perturbations are displayed in Figs. C.1 and C.2 respectively.

### 3.3. Results

The results from testing the proposed mitigation algorithm are summarized in Table 1. Three input perturbations were tested: no-LWE, LWE#1 and LWE#2. The corresponding map of residual wavefront for the different correction basis can be found in Appendix C.

The no-LWE case demonstrates that the proposed mitigation algorithm does not significantly disturb the AO loop. This is already a very important result. More precisely, the small −0.4% SR when the PP control is activated can be ascribed to the suboptimal try-and-error focal-plane analysis. A phase diversity algorithm (for instance F&F) would be required for a more stable, even PP mode measurement.

The LWE#1 perturbation is responsible for −25% SR if no specific LWE correction is applied. Our proposed PTT control allows us to recover +16% SR in a convergence time of about 150 ms (180 loop time steps). Turning on the PP control leads to recovering another +4% SR. The final SR of 85% is only 3.3% behind the best possible correction allowed by the first 990 KL modes.

The LWE#2 perturbation corresponds to a strong LWE that is responsible for a –59% SR loss. Here the PTT control leads to +30% SR and the additional PP control leads to an additional +16% SR. The final SR after the convergence of our mitigation algorithm is only 8% lower than the best achievable correction in

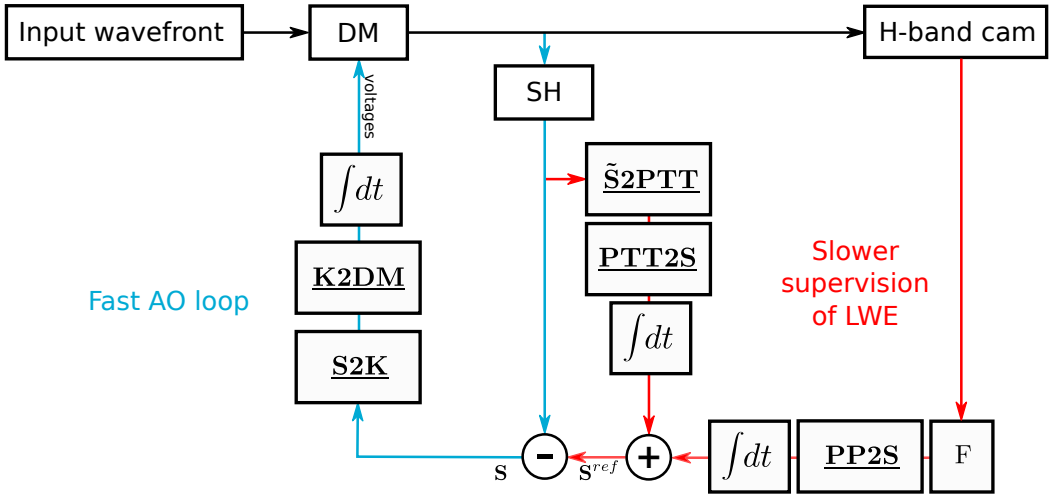


**Fig. 10.** Block diagram for the proposed LWE correction algorithm. The main AO loop is in cyan. The added supervision algorithm is in red.

**Table 1.** Results from the corrective algorithm tests performed in simulation.

| LWE # | Atmos. | LWE rms/ptv (nm) | LWE corr. basis | Final SR | Final rms (nm) |
|---|---|---|---|---|---|
| | Yes | | None | 90.5 ± 2.0% | 81 ± 9 |
| | Yes | | PTT | 90.6 ± 2.0% | 80 ± 9 |
| | Yes | | PTT + PP | 90.1 ± 1.8% | 82 ± 8 |
| 1 | Yes | 173/650 | None | 65 ± 3% | 173 ± 15 |
| 1 | Yes | 173/650 | PTT | 81 ± 5% | 122 ± 18 |
| 1 | Yes | 173/650 | PTT + PP | 85 ± 4% | 106 ± 13 |
| 1 | Yes | 173/650 | Best fitting KL [(a)] | 88.3 ± 2.7% | 92 ± 11 |
| 2 | Yes | 276/1350 | None | 31.2 ± 2.2% | 276 ± 9 |
| 2 | Yes | 276/1350 | PTT | 61 ± 6% | 193 ± 18 |
| 2 | Yes | 276/1350 | PTT + PP | 77 ± 6% | 140 ± 24 |
| 2 | Yes | 276/1350 | Best fitting KL [(a)] | 85.7 ± 2.7% | 107 ± 10 |

**Notes.** In Col. 3, entitled "LWE rms/ptv", values correspond to the post-AO residual without LWE correction. SR values are in *H* band, after convergence of the LWE correction. Values after ± correspond to the standard deviation of the SR and rms on a 10-second sample. [(a)]"Best fitting KL" corresponds to the best possible correction in the 990 KL modes space. To obtain it, we projected the post-AO residuals (without LWE correction algorithm) from the phase space to the KL space. We applied the resulting best-KL fit on an additional DM in our simulated system. The additional DM corrects for the static LWE residuals when the original DM in the AO loop corrects for the atmosphere.

the KL space. We investigated this difference and concluded that our basis (Fig. 8) composed of PTT is not sufficient for controlling the curved content of the LWE aberrations. The algorithm corrects for most of the pupil-scale vortex structure but struggles with the correction of smaller, intricate vortexes around spiders. Deliberately, this LWE#2 perturbation induces strong local vortexes, putting our algorithm in a challenging situation. Still, results show that the proposed mitigation provides a very significant improvement. It also demonstrates that the method has a wide linearity range, allowing it to operate even with strong PTT (1350 nm ptv OPD).

Overall, these results validate the proposed measurement and correction strategies to mitigate the PTT content of the LWE. These results also demonstrate that there is no damaging interplay between the proposed correction algorithm and the atmosphere. The proposed algorithm converges in about 200 AO loop iterations, well in agreement with the low gain 0.005 in the integrator. The effective –3dB correction bandwidth is ≈0.6 Hz (about 100 times slower than the main AO loop).

## 4. Discussion

### *4.1. Advantages and limitations of the mitigation strategy*

The proposed SH-based algorithm dedicated to PTT correction for a partial control of the LWE has some evident advantages. First, the method allows us to recover more than half of the loss of SR due to the LWE, without the use of any additional sensor. Only software modifications are needed to implement it on SPHERE, the AOF or GRAVITY+, for which high-resolution SHs are used. Second, the method is compatible with a subsequent focal-plane analysis in order to control the three PP modes. Even more critically, only one of these three remaining modes is even, and thus requires a fundamentally nonlinear analysis. This simplification is especially interesting when dealing with focal-plane sensing affected by non-common path aberrations. Third, because the method is based on the fast measurement provided by the SH, it has a high temporal bandwidth, sufficient to track LWE temporal evolution (see Sect. 4.2 for a specific discussion on the bandwidth of the LWE). Finally, the method benefits

from the wide linearity range of the SH. In our simulations, the algorithm successfully corrected for PTTs up to 2 µm OPD ptv. This amplitude corresponds to the strongest LWE events documented so far, with less than 10% SR. It is important to notice that the proposed modified controller does not remove the origin of the LWE aberrations (discontinuities), so this notion of capture range does matter.

However, the method obviously suffers from the following limitations. First, it is only a partial solution to tackle the LWE since it only corrects for PTT modes whereas PPs are also significant, low-order contributors. Second, higher-order petal modes (higher than PTTs) are required to correct more complex aberrations introduced by the LWE. Still, their contribution is significantly smaller than the PP and PTT modes (see Appendix D). Third, the correction bandwidths of the turbulence (main AO loop) and of the PTTs (slower modification of the reference slopes) must be sufficiently different to minimize the coupling of the turbulence into the island effect. Indeed, the low orders from the atmosphere, including the TT, project efficiently on the LWE basis. Finally, this mitigation strategy remains in the context of SH spot positioning with CoG, and we have shown that this technique does not provide reliable measurements in LWE conditions (Sect. 2.1). Other positioning techniques such as weighted CoG, thresholding, or match-filtering (Thomas et al. 2006; Ruggiu et al. 1998) might provide a more robust discontinuity measurement and could tackle the LWE problem at the wavefront sensor stage.

### 4.2. Amplitude and bandwidth of the low wind effect

In order to investigate the typical amplitude of the LWE, we analyzed three sequences of ZELDA measurements taken on SPHERE in 2014 (e.g., left panel of Fig. 7). Each sequence was composed of 100 images, sampled at 1.2 s. We projected PPs and PTTs for each pupil-quadrant on the ZELDA phase-screens to estimate quantitatively the contribution of each modes. Combining the three sequences, we obtained the histogram in Fig. E.1 in the appendix. First, this study confirms that the LWE test cases used in the simulation of this paper have typical shape and amplitudes. The linearity of the proposed method is thus largely sufficient to tackle even the worst PTT events. Secondly, this study confirms that PTT modes are important contributors with OPD up to 750 nm ptv. PPs reach at most 400 nm and their distributions have a much narrower range. The proposed method thus provides a significant gain even when restricted to the PTT modes.

We also computed the power spectral density (PSD) of these three ZELDA sequences. We combined the PSDs of all the PP and PTT modes, and of the three ZELDA sequences, in order to improve the overall signal-to-noise ratio. The averaged PSD is displayed in Fig. E.2 in the appendix. The PSD can be approximated by the following model:

$$f < f_c: \quad P(f) \sim (f/f_c)^{-0.4}, \tag{7}$$
$$f > f_c: \quad P(f) \sim (f/f_c)^{-1.3}, \tag{8}$$

with $f$ the frequency and $f_c$ = 0.06 Hz. The knee at the cut-off frequency $f_c$ corresponds to a typical timescale of 16 s. It is well within the ≈0.6 Hz correction bandwidth of the correction algorithm for the PTTs. This validates the requirement that the PTT control loop runs much slower than the main control loop dedicated to atmosphere correction. Moreover, this frequency separation ensures that the gain for the PTTs is kept much smaller than the gain margin of the main AO loop, thus avoiding instabilities.

Both the SPHERE $H$-band differential tip-tilt sensor camera and the GRAVITY $H$-band acquisition camera have frame rates around 1 Hz. Typically, we found that corrections from a focal-plane analysis takes five to ten iterations to converge (gain of 0.1), which corresponds to a -3dB correction bandwidth of ≈0.02 Hz. This is somewhat slower than the LWE cut-off frequency $f_c$ = 0.06 Hz derived above. Again, this highlights the importance of correcting the LWE as much as possible with the information available in the SH.

### 4.3. Understanding whether the low wind effect is global or local

The first LWE study on SPHERE (Sauvage et al. 2015) suggested that the PP and PTT modes were not created by the DM itself, but instead were fully part of the input perturbations, as unseen modes. The present study draws different conclusions. Our simulations show that the AO loop is not blind to PP and PTT. In particular, the simulations explain how a one-quadrant PTT perturbation can lead to a vortex aberration spread over the whole pupil after AO convergence. It indicates that (part of) the LWE problem originates from a faulty response of the AO loop to a peculiar perturbation.

To explore further this aspect, we ran simulations where the input LWE perturbation was not made of PTT- and PP-like modes, but was instead entirely localized along the spiders (see pictures in Appendix F). According to Figs. F.1 and F.2, without the AO loop, those small perturbations lead to a minor decrease in the Strehl ratio (−3 to −6% for a 500 nm OPD ptv perturbation). When closing the loop, the aberration spreads over the whole pupil because of unseen piston and uncorrectable vortex modes, and gives rise to PP- and PTT-like modes. The SR decrease is, therefore, significant (−25 to −45% for a 500 nm OPD ptv perturbation). These basic simulations demonstrate that a perturbation localized close to the spider is sufficient to create the point spread function and post-AO residuals observed on SPHERE, AOF and SCExAO. In fact, when considering the process of spiders cooling the surrounding air, such localized perturbations may appear more realistic than quadrant-scale perturbations. Further studies on AO telemetry data from the instruments affected by the LWE are required to settle the discrepancy.

### 4.4. Pyramid wavefront sensor and discontinuities

Many next-generation instruments will use a pyramid wavefront sensor (e.g., the second stage AO of SPHERE+ and the Single Conjugated AO Natural Guide Star (SCAO NGS) modes of the ELT). In this respect, it is important to discuss if our study on discontinuity measurements by a SH applies to the pyramid wavefront sensor too.

The pyramid wavefront sensor has a two measurement regimes (Vérinaud 2004; Guyon 2005). For low-order modes, the pyramid wavefront sensor measures slopes and has a behavior close to the SH. But for high-order modes, the pyramid behavior tends to an interferometric phase measurement. Vérinaud (2004) shows that the sensor response to a sharp phase step is very different from the SH measurement because the information at the pyramid focal plane is not localized, but is instead distributed on a wide range of subapertures. Our study has shown that, on a SH, only the subapertures directly affected by the discontinuities are (partially) sensitive to the phase step,

resulting in a bad measurement. On a pyramid, the spreading of the signal induced by the discontinuity ensures a good measurement, at least in the weak-phase regime. Also, Bertrou-Cantou et al. (2022) show that, during the best seeing conditions where the LWE occurs, the pyramid wavefront sensor can measure LWE-induced petal modes. However, measurements of phase discontinuities beyond the weak-phase regime still suffer a $\lambda$ phase wrapping. We can conclude that a pyramid in the visible with good seeing conditions (or better, a pyramid in the infrared) is more suitable for the measurement of phase discontinuities than the SH with the classical CoG positioning technique.

*Acknowledgements.* N.P. was supported by the Action Spécifique Haute Résolution Angulaire (ASHRA) of CNRS/INSU co-funded by CNES. The authors acknowledges the support of the French Agence Nationale de la Recherche (ANR), under grant ANR-21-CE31-0017 (project EXOVLTI). The authors would like to thanks Dr. Eric Gendron for pointing us in the direction of curl structures in AO residuals, Dr. Olivier Lai for a very interesting discussion he initiated on LWE mitigation strategies, and Dr. Emiel Por for his kind assistance in the adaptation of HCIPy to our specific problem. We also would like to thanks the GRAVITY+ AO (GPAO) consortium for their support and expertise, Dr. Christophe Vérinaud and Dr. Cédric Taïssir Héritier for their illuminating discussions on pyramids wavefront sensors, and the anonymous referee that helped us to clarify the paper. This research has made use of the following python packages: matplotlib (Hunter 2007), numpy (Harris et al. 2020), hcipy (Por et al. 2018), astropy (Astropy Collaboration 2018) and scipy (Virtanen et al. 2020).

## References

Astropy Collaboration (Price-Whelan, A. M., et al.) 2018, AJ, 156, 123
Baudoz, P., Dorn, R. J., Lizon, J.-L., et al. 2010, SPIE Conf. Ser., 7735, 77355B
Bertrou-Cantou, A., Gendron, E., Rousset, G., et al. 2020, SPIE Conf. Ser., 11448, 1144812
Bertrou-Cantou, A., Gendron, E., Rousset, G., et al. 2022, A&A, 658, A49
Beuzit, J. L., Vigan, A., Mouillet, D., et al. 2019, A&A, 631, A155
Boccaletti, A., Di Folco, E., Pantin, E., et al. 2020, A&A, 637, A5
Bos, S. P., Vievard, S., Wilby, M. J., et al. 2020, A&A, 639, A52
Chanan, G., Troy, M., Dekens, F., et al. 1998, Appl. Opt., 37, 140
Chanan, G., Ohara, C., & Troy, M. 2000, Appl. Opt., 39, 4706
Chauvin, G., Desidera, S., Lagrange, A. M., et al. 2017, A&A, 605, A9
Claudi, R. U., Turatto, M., Gratton, R. G., et al. 2008, SPIE Conf. Ser., 7014, 70143E
Dohlen, K., Langlois, M., Saisse, M., et al. 2008, SPIE Conf. Ser., 7014, 70143L
Fried, D. L. 1998, J. Opt. Soc. Am. A, 15, 2759
Fusco, T., Sauvage, J. F., Mouillet, D., et al. 2016, SPIE Conf. Ser., 9909, 99090U
Ginski, C., Facchini, S., Huang, J., et al. 2021, ApJ, 908, L25
Gonsalves, R. A. 2002, in European Southern Observatory Conference and Workshop Proceedings, 58, 121
Guyon, O. 2005, ApJ, 629, 592
Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Nature, 585, 357
Holzlöhner, R., Kimeswenger, S., Kausch, W., & Noll, S. 2021, A&A, 645, A32
Hunter, J. D. 2007, Comput. Sci. Eng., 9, 90
Hutterer, V., Shatokhina, I., Obereder, A., & Ramlau, R. 2018, J. Astron. Telescopes Instrum. Syst., 4, 049005
Keller, C. U., Korkiakoski, V., Doelman, N., et al. 2012, SPIE Conf. Ser., 8447, 844721
Keppler, M., Benisty, M., Müller, A., et al. 2018, A&A, 617, A44
Korkiakoski, V., Keller, C. U., Doelman, N., et al. 2012, SPIE Conf. Ser., 8447, 84475Z
Korkiakoski, V., Keller, C. U., Doelman, N., et al. 2014, Appl. Opt., 53, 4565
Langlois, M., Gratton, R., Lagrange, A. M., et al. 2021, A&A, 651, A71
Luo, J., Huang, H., Matsui, Y., et al. 2015, Opt. Express, 23, 8706
Milli, J., Vigan, A., Mouillet, D., et al. 2017, A&A, 599, A108
Milli, J., Kasper, M., Bourget, P., et al. 2018, SPIE Conf. Ser., 10703, 107032A
Mouillet, D., Milli, J., Sauvage, J. F., et al. 2018, SPIE Conf. Ser., 10703, 107031Q
N'Diaye, M., Dohlen, K., Fusco, T., & Paul, B. 2013, A&A, 555, A94
N'Diaye, M., Vigan, A., Dohlen, K., et al. 2016, A&A, 592, A79
N'Diaye, M., Martinache, F., Jovanovic, N., et al. 2018, A&A, 610, A18
Oberti, S., Kolb, J., Madec, P.-Y., et al. 2018, SPIE Conf. Ser., 10703, 107031G
Por, E. H., Haffert, S. Y., Radhakrishnan, V. M., et al. 2018, in Proc. SPIE, Vol. 10703, Adaptive Optics Systems VI
Primmerman, C. A., Price, T. R., Humphreys, R. A., et al. 1995, Appl. Opt., 34, 2081
Ruggiu, J.-M., Solomon, C. J., & Loos, G. 1998, Opt. Lett., 23, 235
Sauvage, J.-F., Fusco, T., Guesalaga, A., et al. 2015, in Adaptive Optics for Extremely Large Telescopes IV (AO4ELT4), E9
Sauvage, J.-F., Fusco, T., Lamb, M., et al. 2016a, SPIE Conf. Ser., 9909, 990916
Sauvage, J.-F., Fusco, T., Petit, C., et al. 2016b, J. Astron. Telescopes Instrum. Syst., 2, 025003
Schmid, H. M., Bazzon, A., Roelfsema, R., et al. 2018, A&A, 619, A9
Schwartz, N., Sauvage, J.-F., Correia, C., et al. 2017, AO4ELT5 Proceedings, https://doi.org/10.26698/AO4ELT5.0015
Thomas, S., Fusco, T., Tokovinin, A., et al. 2006, MNRAS, 371, 323
Tyler, G. A. 2000, J. Opt. Soc. Am. A, 17, 1828
van Boekel, R., Henning, T., Menu, J., et al. 2017, ApJ, 837, 132
van Dam, M. A., Raglandb, S., & Wizinowichb, P. L. 2017, AO4ELT5 Proceedings
Vérinaud, C. 2004, Opt. Commun., 233, 27
Vievard, S., Bos, S., Cassaing, F., et al. 2019, ArXiv e-prints, [arXiv:1912.10179]
Vigan, A., Fontanive, C., Meyer, M., et al. 2021, A&A, 651, A72
Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nat. Methods, 17, 261
Wilby, M. J., Keller, C. U., Sauvage, J. F., et al. 2016, SPIE Conf. Ser., 9909, 99096C
Wilby, M. J., Keller, C. U., Sauvage, J. F., et al. 2018, A&A, 615, A34
Wu, T., Berto, P., & Guillon, M. 2021, Appl. Phys. Lett., 118, 251102

## Appendix A: Adaptive optics residuals with respect to the position of the discontinuity

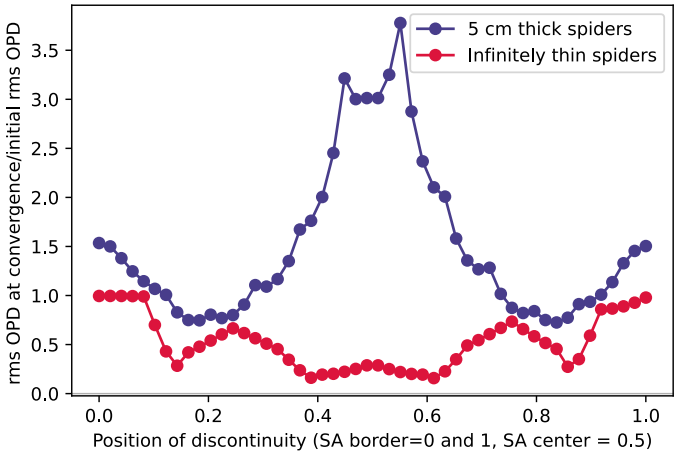


Fig. A.1: Ratio between the rms of the input aberration and the rms of the residual wavefront at convergence, when varying the position of the spiders with respect to the subapertures. The AO simulation is described in Sect. 2.2.1. The input aberration is the 80 nm OPD ptv PTT displayed in Fig. 5. The residual wavefront detailed in Fig. 5 corresponds to the leftmost red point (position = 0, infinitely thin spiders).

## Appendix B: Derivations for the focal-plane analysis

In order to estimate the PP modes from the focal-plane images, we implemented a simple algorithm based on F&F, the derivations of which can be found in Korkiakoski et al. (2014). In our case the focal-plane analysis is restricted to the three PPs. We analytically derived the Fraunhofer approximation (Eq. (B.1)) in the weak-phase regime (Eq. (B.2)), thus limiting to first order:

$$\mathbf{I} = |\,\mathscr{F}\{\mathbf{A}\exp(i\boldsymbol{\phi})\}\,|^2, \tag{B.1}$$

$$\mathbf{I} \approx |\,\mathscr{F}\{\mathbf{A} + i\mathbf{A}\boldsymbol{\phi})\}\,|^2, \tag{B.2}$$

where $\mathbf{I}$ is the image, $\mathscr{F}$ is the Fourier transform operation, $\mathbf{A}$ is the pupil transmission, and $\boldsymbol{\phi}$ is the phase screen. Reducing to the three PP modes of Fig. 8 we get:

$$\boldsymbol{\phi} = P_1\,\boldsymbol{\phi}_1 + P_2\,\boldsymbol{\phi}_2 + P_3\,\boldsymbol{\phi}_3,$$

where $\boldsymbol{\phi}_1$, $\boldsymbol{\phi}_2$, and $\boldsymbol{\phi}_3$ are the normalized wavefronts of the three PP modes, and $P_1$, $P_2$, and $P_3$ are the corresponding amplitudes of these modes. Then,

$$\mathbf{I} \approx |\,\mathscr{F}\{\mathbf{A} + i\,P_1\mathbf{A}\boldsymbol{\phi}_1 + i\,P_2\mathbf{A}\boldsymbol{\phi}_2 + i\,P_3\mathbf{A}\boldsymbol{\phi}_3)\}\,|^2.$$

We define the following Fourier Transforms:

$$\tilde{\mathbf{A_0}} = \mathscr{F}\{\mathbf{A}\},$$
$$\tilde{\mathbf{A_1}} = \mathrm{Im}\{\,\mathscr{F}\{\mathbf{A}\boldsymbol{\phi}_1\}\,\},$$
$$\tilde{\mathbf{A_2}} = \mathscr{F}\{\mathbf{A}\boldsymbol{\phi}_2\},$$
$$\tilde{\mathbf{A_3}} = \mathrm{Im}\{\,\mathscr{F}\{\mathbf{A}\boldsymbol{\phi}_3\}\,\}.$$

The focal plane image $\mathbf{I}$ can be decomposed into the sum of an even image $\mathbf{I}_e$ and an odd image $\mathbf{I}_o$:

$$\mathbf{I} = \mathbf{I}_e + \mathbf{I}_o. \tag{B.3}$$

By investigating the symmetry of the various terms, we obtain:

$$\mathbf{I}_e \approx \tilde{\mathbf{A_0}}^2 + P_1^2\tilde{\mathbf{A_1}}^2 + P_2^2\tilde{\mathbf{A_2}}^2 + P_3^2\tilde{\mathbf{A_3}}^2 + 2\,P_1\,P_3\,\tilde{\mathbf{A_1}}\,\tilde{\mathbf{A_3}}, \tag{B.4}$$

$$\mathbf{I}_o \approx -2\,P_1\,\tilde{\mathbf{A_0}}\tilde{\mathbf{A_1}} - 2\,P_3\,\tilde{\mathbf{A_0}}\tilde{\mathbf{A_3}}. \tag{B.5}$$

Equation (B.5) defines a linear transform between the unknowns $P_1$ and $P_3$, and the observable $\mathbf{I}_o$. The set of two vectors $-2\tilde{\mathbf{A_0}}\tilde{\mathbf{A_1}}$ and $-2\tilde{\mathbf{A_0}}\tilde{\mathbf{A_3}}$ defines an interaction matrix, which can be calibrated by playing known perturbations. With the inverted interaction matrix, we obtain amplitudes of the two odd PP modes for each frame.

Equation (B.4) is nonlinear and has a fundamental sign ambiguity. To deal with the nonlinearity, we rearrange Eq. (B.5) and Eq. (B.4) into:

$$P_2\tilde{\mathbf{A_2}} \approx \sqrt{\mathbf{I}'_e - \tilde{\mathbf{A_0}}^2 - \frac{\mathbf{I}_o^2}{4\tilde{\mathbf{A_0}}^2}}, \tag{B.6}$$

where we applied the necessary normalization on $\mathbf{I}_e$ (described in Korkiakoski et al. (2014)):

$$\mathbf{I}'_e = \mathbf{I}_e + \left(1 - \frac{\max(\mathbf{I}_e)}{\max(\tilde{\mathbf{A_0}}^2)}\right)\tilde{\mathbf{A_0}}^2. \tag{B.7}$$

We found that the median of the right-hand side of Eq. (B.6) is a robust estimator of $P_2$. This estimate is multiplied by an unknown, constant scalar factor (the median of $\tilde{\mathbf{A_2}}$), which is calibrated manually. To lift the sign ambiguity, we implemented a (suboptimal) try-and-error approach.

## Appendix C: Wavefronts corrected by the mitigation algorithm

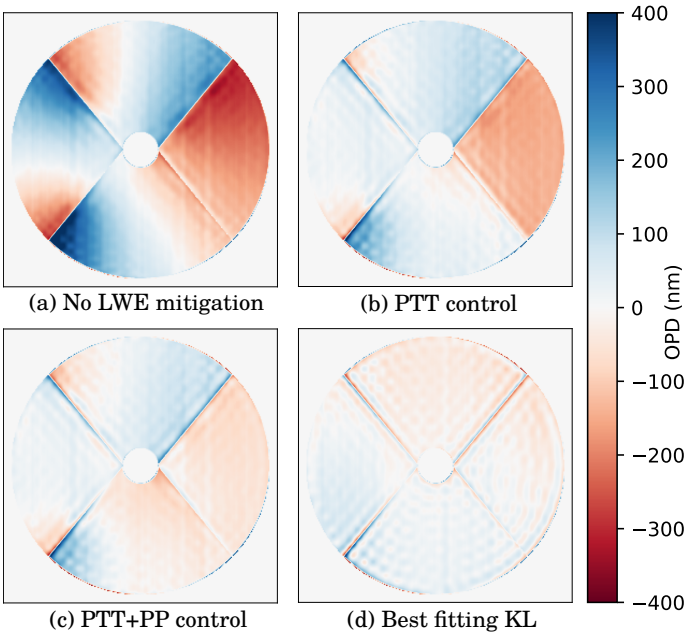


(a) No LWE mitigation (b) PTT control

(c) PTT+PP control (d) Best fitting KL

Fig. C.1: Post-AO residuals for the LWE#1 perturbation. Each (a), (b), (c), and (d) wavefront shows the OPD at the mitigation algorithm convergence for the corresponding line in Table 1. Wavefronts are averaged over 2 s to wash out the atmosphere contribution.

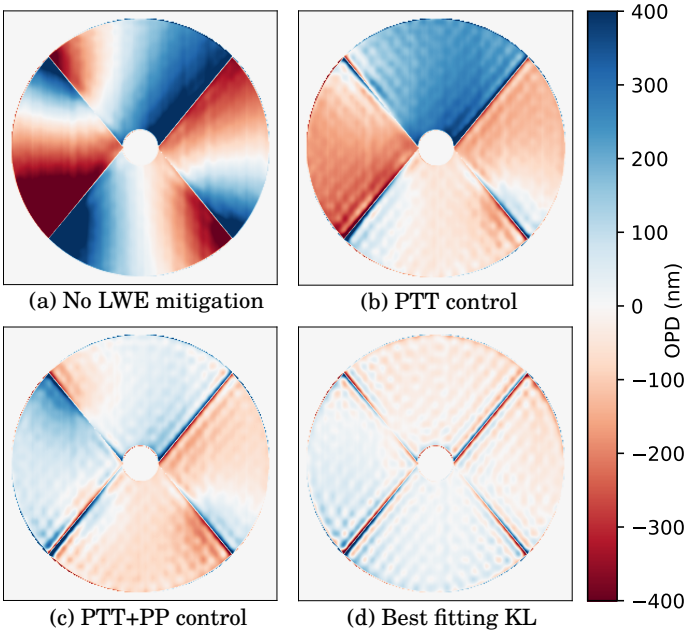


(a) No LWE mitigation (b) PTT control

(c) PTT+PP control (d) Best fitting KL

Fig. C.2: Post-AO residuals for the LWE#2 perturbation. Each (a), (b), (c), and (d) wavefront shows the OPD at the mitigation algorithm convergence for the corresponding line in Table 1. Wavefronts are averaged over 2 s to wash out the atmosphere contribution.

## Appendix D: Power in the petaling modes higher than petal-tip-tilts

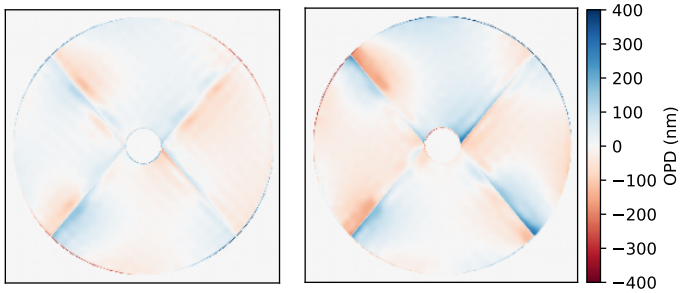


(a) LWE#1 (b) LWE#2

Fig. D.1: Post-AO residuals for LWE#1 and LWE#2 when assuming an ideal correction of PP and PTT. All that remains is the higher-order content composed of vortexes around the spiders.

**Table D.1.** Comparison of the post-AO residuals for the LWE#1 and LWE#2 perturbations, with and without an ideal PP and PTT correction.

| LWE # | | SR | rms (nm) |
|---|---|---|---|
| 1 | Uncorrected post-AO residuals | 70 %$^a$ | 160$^a$ |
| 1 | PP+PTT ideal correction | 98 %$^b$ | 37$^b$ |
| 2 | Uncorrected post-AO residuals | 34 %$^c$ | 263$^c$ |
| 2 | PP+PTT ideal correction | 96 %$^d$ | 55$^d$ |

**Notes.** Residuals after the ideal PP+PTT correction give the power of the higher-order LWE petaling modes. These values are taken without atmosphere. SR values are in H band. $^{(a)}$ corresponds to Fig. C.1.a but without the atmosphere contribution. $^{(b)}$ corresponds to the OPD screen on Fig. D.1a. $^{(c)}$ corresponds to Fig. C.2.a but without the atmosphere contribution. $^{(d)}$ corresponds to the OPD screen on Fig. D.1b.

## Appendix E: Analysis of ZELDA low-wind-effect sequences

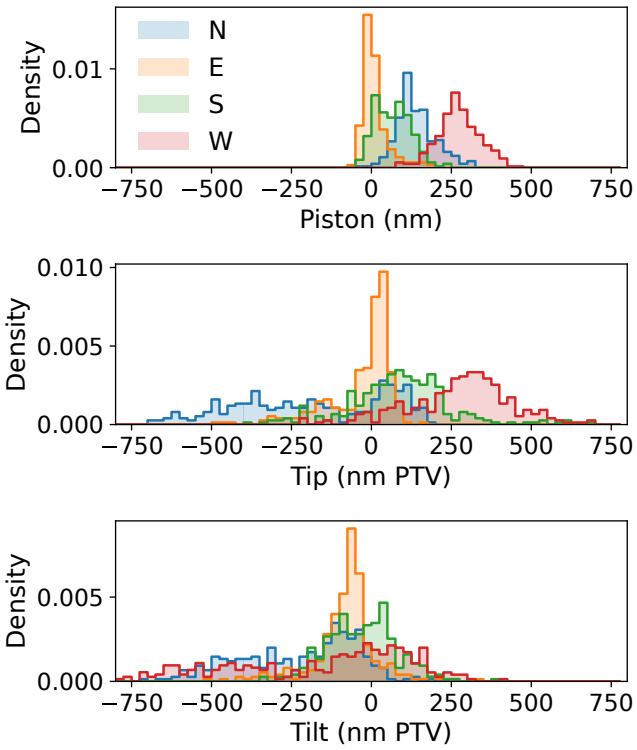


Fig. E.1: Distribution of piston, tip and tilt per quadrant (north, east, south and west) for a combination of three ZELDA sequences measured on SPHERE during the LWE. The piston reference is arbitrarily set to be 0 on average on the east quadrant.

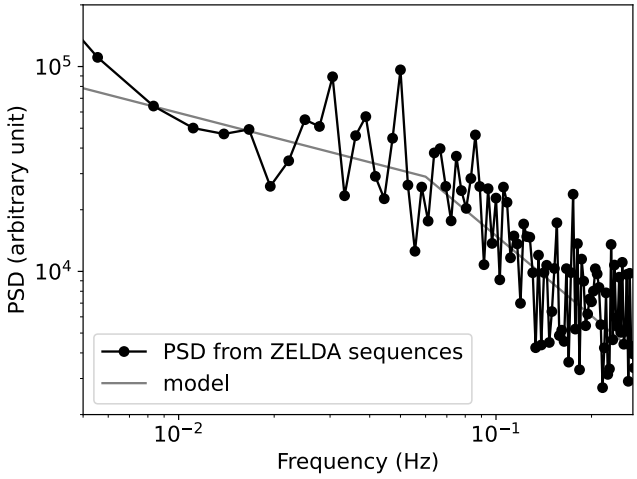


Fig. E.2: Power spectral density calculated from three sequences of wavefront measurements with the ZELDA sensor, during LWE events on SPHERE. ZELDA measurements are projected on a PP and PTT basis, and the power spectra of each mode is summed. The PSD model displays a knee at $f_c = 0.06$ Hz.

## Appendix F: Adaptive optics response to local low wind effect

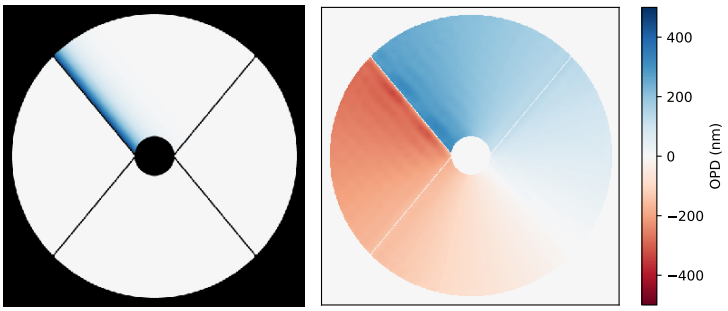


(a) Perturbation (b) Post-AO residuals

Fig. F.1: AO response to an OPD gradient localized along a spider with 500 nm OPD ptv. The OPD has a discontinuity at the position of the spider, and decreases exponentially with the distance to the spider. There is no atmosphere. The SR at H-band for the perturbation before closing the AO loop is 97% (a). The SR at AO convergence is 75% (b).

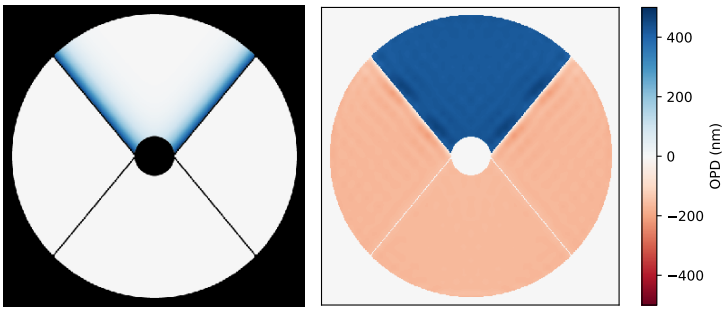


(a) Perturbation (b) Post-AO residuals

Fig. F.2: AO response to two OPD gradients localized along different spiders with 500 nm OPD ptv each. The OPD has a discontinuity at the position of the spider, and decreases exponentially with the distance to the spider. There is no atmosphere. The SR at H-band for the perturbation before closing the AO loop is 94% (a). The SR at AO convergence is 55% (b).