



Universiteit
Leiden
The Netherlands

Painting with starlight : optical techniques for the high-contrast imaging of exoplanets

Wilby, M.J.

Citation

Wilby, M. J. (2018, November 27). *Painting with starlight : optical techniques for the high-contrast imaging of exoplanets*. Retrieved from <https://hdl.handle.net/1887/67531>

Version: Not Applicable (or Unknown)

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/67531>

Note: To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The handle <http://hdl.handle.net/1887/67531> holds various files of this Leiden University dissertation.

Author: Wilby, M.J.

Title: Painting with starlight : optical techniques for the high-contrast imaging of exoplanets

Issue Date: 2018-11-27

Chapter 3

A ‘Fast & Furious’ solution to the low-wind effect for SPHERE at the VLT

Michael J. Wilby, Christoph U. Keller, Jean-François Sauvage, Thierry Fusco, David Mouillet, Jean-Luc Beuzit, and Kjetil Dohlen
Proc. SPIE 9909, Adaptive Optics Systems V, 99096C (2016)

Abstract

We present a potential non-invasive solution to sensing the so-called low-wind effect (LWE) seen in the SPHERE instrument at the VLT, based on the ‘Fast and Furious’ (F&F) sequential phase diversity wavefront reconstruction algorithm. This uses non-coronagraphic focal-plane images available from the near-infrared differential tip-tilt sensor (DTTS), with the closed-loop correction cycle itself providing the necessary phase diversity between frames required to reconstruct the full wavefront phase. Crucially, this means F&F does not need to apply large artificial phase probes as required by standard phase diversity algorithms, allowing it to operate in a real-time (~ 10 Hz) correction mode without impacting science observations. In this paper we present the results of realistic closed-loop AO simulations designed to emulate SPHERE/DTTS observations of the LWE. With this we demonstrate that the F&F algorithm is capable of effective removal of the characteristic point-spread function (PSF) aberrations of strongly LWE-affected images within a few closed-loop iterations, with the final wavefront quality limited only by the corrective order of the deformable mirror. The ultimate goal of this project is to provide an independent, real-time and focal-plane wavefront sensor for SPHERE which is capable of detecting and directly compensating the LWE as it arises, thus improving coronagraph performance under the best 15-20% of observing conditions where the effect is most pronounced.

3.1 Introduction

3.1.1 The low-wind effect

The Spectro-Polarimetric High-contrast Exoplanet REsearch instrument (SPHERE, Beuzit et al. 2008) for the Very Large Telescope (VLT) has recently finished its commissioning phase and is already providing exceptional results in the field of direct exoplanet and circumstellar disk imaging (Vigan et al. 2016a; Maire et al. 2016; Zurlo et al. 2016; Bonnefoy et al. 2016). The performance of the instrument is however limited under the best seeing conditions by the so-called low-wind effect (LWE), where the un-occulted stellar point spread function (PSF) gains two or more bright side-lobes (dubbed ‘Mickey Mouse ears’) on spatial scales of 2 to 4 λ/D and relative intensity of 10^{-1} or greater with respect to the central PSF core; a number of examples of the effect are shown in Fig 3.1 as observed by the SPHERE differential tip-tilt sensor (DTTS). The appearance of the LWE is strongly correlated with the most favourable observing conditions where the dome-level wind speed at the VLT drops below 3 ms^{-1} . The effect is present at some level in 15-20% of observations and therefore leads to significant degradation of the performance of the focal-plane coronagraph masks of both the visible and near-infrared (NIR) instrument arms of SPHERE.

In the current absence of direct sensing or control solutions for the LWE, the inclusion of an additional and complementary wavefront sensor is necessary in order to identify and suppress this effect without interrupting nominal science observations. The prototype Zernike sensor for Extremely Low-level Differential Aberrations (ZELDA, N’Diaye et al. 2014; N’Diaye et al. 2016) was installed into the NIR coronagraph filter wheel in 2014, converting the science camera of the IRDIS subsystem into a pupil-plane wavefront sensor. The high-resolution phase maps obtained during its testing phase provided a clear interpretation of the observed LWE-affected PSFs of Fig. 3.1, and form the basis for the wavefront models used in this paper. However, before a proposed upgrade to SPHERE is performed in which ZELDA can be fully integrated as a stand-alone wavefront sensor, it is currently not possible to operate the sensor in parallel with IRDIS or IFS science observations. An example wavefront model and corresponding PSF based on ZELDA data is shown in Fig 3.2. From this it can be seen that each ‘ear’ present in the PSF is a direct result of strong differential piston, tip and tilt phase error components over a single VLT pupil segment, resulting in quadrant-to-quadrant phase discontinuities which have been observed on-sky to reach amplitudes of up to 800 nm.

The current most widely accepted hypothesis regarding the origin of the LWE attributes it to changing internal dome conditions, whereby slow wind flowing across the VLT secondary mirror support spiders allows for significant thermal transfer between air and the metal surface, creating an optical path difference across each spider which gradually reduces in amplitude further downwind as the air returns to normal temperature. However, the resulting wavefront is entirely invisible to the filtered Shack-Hartmann wavefront sensor of the SPHERE extreme adaptive optics (XAO) system, SAXO (Sauvage et al. 2014; Fusco et al. 2014), making it difficult to explain the existence of quadrant-sized tip/tilt errors using only sources upstream of SPHERE. Attempts to eliminate the effect via modification of SAXO closed-loop parameters or the management of dome airflow conditions have thus far been unable to demonstrably reduce the strength of effect (J.-F. Sauvage & J.-L. Beuzit, private communication,




Figure 3.1: Example images of the diverse PSF morphologies which have been associated with the LWE, obtained with the SPHERE DTTS. The most extreme case (lower right) corresponds to significant tip/tilt error across all 4 VLT pupil segments and a central Strehl ratio of approximately zero. (Adapted from the SPHERE User Manual v.96, ESO 2015)




Figure 3.2: *Left:* Model VLT pupil with an example LWE wavefront based on measurements made with the ZELDA wavefront sensor, including phase error model due to faulty deformable mirror actuators. *Right:* Corresponding aberrated PSF displaying two main diffracted lobes at the location of the first airy ring ($\sim 2.5\lambda/D$), which correspond directly with tip/tilt error across individual pupil segments in the phase error map.

2016).

3.1.2 Proposal: Focal-plane wavefront sensing using the SPHERE DTTS

In the absence of a dedicated LWE wavefront sensor within the commissioned SPHERE instrument, it is necessary in the short-term to develop a wavefront sensing approach which simultaneously meets the following criteria:

- 1) Provides continuous estimates of the LWE wavefront phase on second-to-minute timescales, matching the fastest observed evolution of the effect.
- 2) Uses only hardware currently implemented in SPHERE
- 3) Does not generate additional overhead or down-time for science observations
- 4) Compensates measured LWE without impeding the main atmospheric correction loop of SAXO.

In this paper we propose the so-called Fast and Furious (F&F) sequential phase diversity algorithm (Keller et al. 2012; Korkiakoski et al. 2014) as a focal-plane wavefront sensor which meets all of the above criteria. In order to work effectively this algorithm uses sequential images of the narrowband, non-coronagraphic stellar PSF, with frames separated by known (but not manually controlled) phase offsets. Unlike standard phase diversity, which typically requires one image in each pair to be strongly aberrated by a known phase probe such as defocus (Gonsalves 1982; Sauvage et al. 2007), F&F uses the smaller frame-to-frame phase diversity generated by actively correcting the measured aberrations to provide the additional information necessary to reconstruct the full wavefront phase.

Fortunately a subsystem already exists within the SPHERE instrument which meets these requirements, in the form of the DTTS imaging camera (Baudoz et al. 2010). This consists of a Hawaii I 1024×1024 detector operating at H-band ($\lambda_0 = 1.53 \mu\text{m}$, 3% bandwidth) located behind a 2% beam splitter immediately upstream of the NIR coronagraph wheel, as shown in Fig. 3.3. During normal observations this runs at frame-rates between 1-1000 Hz (depending on target magnitude) to ensure the stellar image remains accurately centred on the focal-plane mask of the NIR Lyot coronagraph. These properties match extremely well with the requirements of F&F, making the DTTS a good source of real-time focal-plane images for use with the algorithm.

F&F has already been shown to be capable of performing reliable closed-loop PSF optimisation in idealised simulations (Keller et al. 2012) and optical bench tests (Korkiakoski et al. 2014) for the purpose of sensing and correcting non-common-path wavefront errors (NCPEs). The on-sky performance of the algorithm or its behaviour in the presence of complicating factors, such as residual incoherent atmospheric speckles and significant image noise, has however not yet been studied. The main goal of this paper is therefore to verify that the F&F algorithm is also able to robustly compensate the LWE under realistic SPHERE observing conditions. To this end, in this paper we present the results of closed-loop simulations which emulate real DTTS images by including factors such as the strong LWE, XAO-corrected atmospheric speckles, NCPEs and significant detector noise sources.




Figure 3.3: Block diagram of the SPHERE instrument, indicating the location of the DTTS directly upstream of the NIR coronagraph focal plane. (Adapted from Beuzit et al. 2008)

3.2 The Fast & Furious algorithm

The term ‘Fast and Furious’ here refers to a modified sequential phase diversity algorithm developed in Keller et al. (2012) and Korkiakoski et al. (2014), designed for focal-plane phase-retrieval without requiring the application of PSF-degrading phase probes. This makes it suitable to sense NCPEs in real time using the science focal plane, without interrupting science observations. F&F achieves this by using a second-order expansion of the aberrated stellar intensity image in terms of the un-aberrated telescope PSF and the Fourier transforms of the odd and even aberrating wavefront for which we need to solve. This solution is degenerate in the signs of the even focal-plane field, which must be broken with the use of a phase diversity; this is provided by the previous image in the observation sequence, for which the wavefront differs by a known phase offset.

This approach has the following advantages over conventional phase diversity and other focal-plane phase retrieval algorithms, such as the differential optical transfer function (dOTF) approach (Codona 2013):

- 1) No destructive phase probes or amplitude masks: closed-loop update cycle provides the necessary PSF diversity
- 2) Fast: requires only one 2D spatial Fourier transform and a small number of linear operations per iteration
- 3) Requires no physical modification to instrument components during operation (e.g. camera defocus)
- 4) Makes no assumptions on the spatial order or specific morphology of the wavefront phase.

Below we present a summary of the key equations required for a single closed-loop iteration of F&F, the full details of which can be found in Korkiakoski et al. (2014).

This process is also illustrated schematically in Fig. 3.4 for an example wavefront containing one even and one odd Zernike mode. For convenience, focal-plane quantities are denoted by lower case variables, while the corresponding pupil-plane quantities are represented by the equivalent upper-case variables. All quantities are two-dimensional unless otherwise specified.

The second-order expansion of the aberrated PSF can be written as

$$p = S a^2 + 2a(ia * \phi_o) + (ia * \phi_o)^2 + (a * \phi_e)^2 \quad (3.1)$$

$$= S a^2 + 2ay \quad + y^2 \quad + v^2, \quad (3.2)$$

in terms of the complex un-aberrated focal-plane electric field $a = \mathcal{F}[A]$, where A is the telescope aperture function, and the even and odd components of the aberrating focal-plane fields $\phi_{o,e} = \mathcal{F}[\Phi_{o,e}]$ respectively, with odd/even symmetry defined to be consistent with that of the Fourier transform operator $\mathcal{F}[\cdot]$. $*$ denotes the convolution operator. The scalar normalisation factor $S \approx (1 - \sigma_\phi^2)$ is the image Strehl ratio, applied as a second-order correction to improve energy conservation, where σ_ϕ^2 is the wavefront variance. For convenience we have defined the following focal plane quantities y, v as the (complex) focal-plane fields of the even and odd aberrating wavefronts, following the notation of Gonsalves (2001):

$$y = i\mathcal{F}[A\Phi_o] = (ia * \phi_o), \text{ and} \quad (3.3)$$

$$v = \mathcal{F}[A\Phi_e] = (a * \phi_e), \quad (3.4)$$

which for an even aperture function are real quantities due to Fourier transform symmetries, although should in general be treated as complex for real aperture functions such as the VLT pupil which include asymmetric spiders. The problem then reduces to estimating v and y by splitting Eq. 3.2 into its even and odd components and solving, yielding

$$y = ap_o/(2a^2 + \epsilon), \text{ and} \quad (3.5)$$

$$|v| = \sqrt{|p_e - (S a^2 + y^2)|}. \quad (3.6)$$

Here ϵ is a scalar noise reduction term, typically chosen to be one order of magnitude greater than the noise background in the image frames. This solution is degenerate in sign of v , where the sign ambiguity associated with the square root must be broken with the use of the second focal-plane image of known phase diversity. The map of signs of v is then computed by solving the even part of the two-equation system consisting of Eq. 3.2 and its analogue for the previous image p_1 , which differs from the current image p_2 by a known phase diversity Φ_d :

$$\text{sign}(v) = \text{sign}\left(\frac{p_{2,e} - p_{1,e} - (v_d^2 + y_d^2 + 2yy_d)}{2v_d}\right), \quad (3.7)$$

where $v_d = \mathcal{F}[A\Phi_{d,e}]$, $y_d = i\mathcal{F}[A\Phi_{d,o}]$ are the Fourier transforms of the even/odd components of the known phase diversity Φ_d separating the chosen PSF images p_1 and p_2 . This estimator is not used to calculate the pixel amplitudes of v because it involves the subtraction of two similar PSFs, resulting in elevated pixel-to-pixel noise with respect to Eq. 3.6. Because the first iteration of F&F must be performed without access




Figure 3.4: Schematic illustration of the F&F algorithm detailed in Eq. 3.5-3.8, showing the second iteration of the algorithm (i.e. the first iteration using full phase diversity information) for an example wavefront initially containing 0.25 radians root-mean-square (RMS) phase error of each of astigmatism (Z_5) and coma (Z_6) Zernike modes. The most recent PSF p_2 provides a full estimate of the odd wavefront phase Φ_o and the absolute value of the even focal-plane field $|v|$. Using p_1 and Φ_d from the previous iteration provides an estimate of the sign of v and hence completes the wavefront phase reconstruction Φ_2 . This is then passed to the deformable element and becomes the phase diversity for the next iteration step.

to phase diversity information, it is convenient to take the signs of the un-aberrated field, i.e. $\text{sign}(v) = \text{sign}(a)$, for this step only. This is found in practice to improve convergence speed over the case where all signs are assumed to be positive, and robustly results in a decrease in overall wavefront variance in the first iteration.

The final wavefront estimated is then computed in an efficient manner with the single inverse Fourier transform

$$A\Phi = \mathcal{F}^{-1}[\text{sign}(v)|v| - iy], \quad (3.8)$$

where the output may be spatially filtered or projected onto the desired correcting mode basis to reduce pixel-to-pixel noise in the final wavefront estimate.

3.3 Simulating F&F performance in SPHERE

3.3.1 Simulation parameters

The goal of this paper is to demonstrate that the F&F algorithm is capable of reliably eliminating LWE-like wavefronts given a noisy image sequence supplied by the SPHERE DTTS. The following parameters are therefore included in these simulations:

- 1) **Low-wind effect:** Uses the LWE phase screen shown in Fig. 3.2, which is allowed to grow linearly in amplitude from zero to a maximum discontinuity of 800 nm across the VLT spiders (equal to 1200 nm peak-valley error), over a timescale of 10 seconds.
- 2) **XAO-corrected atmospheric speckles:** Generated using a Kolmogorov spatial power spectrum with three wind layers of effective $r_0 = 30$ cm, at an observing wavelength of $1.5 \mu\text{m}$. A spatial filter is applied to low frequencies to approximate a SAXO-like control radius of approximately $13 \lambda/D$ (see top panel of Fig. 3.5). The filtered phase screens are passed over the aperture with a coherence time-scale of $\tau_0 \approx 1$ ms, producing a mean Strehl ratio of $S \approx 0.75$.
- 3) **Non-common-path errors:** Aberrations are generated by randomly drawing coefficients from a 100-mode Zernike basis with a simple flat spatial power spectral density. The coefficients are individually convolved with Gaussian temporal evolution profiles with a mean full width at half maximum of 1 s and the resulting time-series normalised to provide a mean RMS wavefront error of 75 nm. An example NCPE coefficient time series and wavefront are shown in the lower panel of Fig. 3.5.
- 4) **Detector Noise:** The simulated 512×512 pixel detector has a resolution of 3 pixels per λ/D , producing simulated DTTS-like images with an effective exposure time of 0.1 s. This includes a constant intensity background at 5×10^{-3} contrast and including photon noise, $50 \text{ e}^-/\text{s}$ dark current and 20 e^- readout noise sources with a final 16-bit analog-to-digital unit (ADU) conversion matched to an effective well depth of 100,000 e^- and bias offset of 500 e^- .
- 5) **Corrective element:** The final wavefront estimates are spatially filtered by projection onto a basis of the first 50 Zernike modes, as a conservative estimate of the corrective ability of the 41×41 actuator SAXO deformable mirror (DM). This filtered correction map is then assumed to be implemented accurately on




Figure 3.5: *Top:* XAO-filtered spatial power spectrum used to generate atmospheric speckle phase screens, with the x-scale converted to units of λ/D to correspond directly with focal-plane coordinates. The dashed line indicates the $k^{-11/3}$ power law of Kolmogorov turbulence. *Bottom:* Total non-common-path RMS wavefront error (top) and three examples of the temporal evolution of single Zernike mode coefficients (bottom), with a coherence time of ~ 1 s. Inset: Representative NCPE phase map comprising all 100 Zernike modes, evaluated at $t = 7$ s.

the DM at a cadence of 10 Hz to match the simulated DTTS exposures, with zero frames of lag.

In order to test the limiting performance of the algorithm, these parameters have been deliberately adjusted to provide a conservative estimate of SPHERE performance: the correction radius, H-band Strehl ratio and residual speckle noise inside the control radius have all been deliberately degraded with respect to reported SAXO performance (Fusco et al. 2014; Sauvage et al. 2015). The aim of this is to maximise the potential impact of noise sources (in particular the presence of an incoherent speckle background) on F&F, such that a successful verification in simulation provides greater confidence in the on-sky performance of the algorithm under less strenuous conditions. Examples of the simulated DTTS images generated using the above parameters are shown in Fig. 3.6. Here the control radius due to the XAO power spectrum can be seen at $\sim 13\lambda/D$, although this and the majority of diffracted intensity structure beyond the first Airy ring are hidden beneath the detector noise floor in final images. This is comparable to the observed image quality presented in Fig. 3.1 whilst still containing all the major factors against which the F&F algorithm must be tested in order to evaluate its potential for on-sky operation.

3.3.2 Simulation results

Figure 3.7 shows the results of a demonstrative closed-loop F&F simulation with the parameters specified in the previous section. The total duration of 30 seconds spans $3 \times 10^4 \tau_0$ and 30 NCPE coherence time-scales, and as such fully samples the behaviour of the secondary aberration sources. The first 5 seconds involve only these two components, after which a LWE is allowed to slowly evolve over the course of 10 seconds, reaching the maximum 800 nm phase discontinuity amplitude at the 15 s half-way point in the simulation. In order to simulate a scenario in which F&F is activated once the LWE reaches a critical aberration threshold, correction begins at the 10 second mark (blue curve), when the LWE amplitude reaches 400 nm. The final 15 seconds after the LWE has finished evolving are then used to produce long integrations of the simulated post-correction DTTS images shown in the left panels of Fig. 3.7, which are used to evaluate the performance of the algorithm with respect to the un-corrected case.

The final LWE discontinuity amplitude of 800 nm represents an extreme case exceeding the maximum effect observed on-sky, resulting in a final Strehl of less than 15% once residual atmospheric speckles and NCPEs are included. Nonetheless, it can be seen that the F&F algorithm is able to effectively remove the characteristic ‘Mickey Mouse ears’ by compensating only the first 50 lowest-order Zernike modes, recovering a diffraction limited core and first Airy ring along with a significant gain in image Strehl ratio. The convergence of the algorithm is also rapid once initialised; for a closed-loop gain of 0.3 the Strehl ratio increases by 20% within the first five iterations. It can be seen however that image quality then continues to degrade slowly up to 15 s simulation time as the strength of the LWE continues to increase, with the final converged wavefront Strehl ratio stabilising at approximately 60%, 16% below its initial pre-LWE value. This loss of PSF quality can be explained entirely by the high-order components of the injected LWE wavefront which are not addressed when correcting only the first 50 Zernike modes. This is shown in Fig. 3.8, where the majority of the residual aberration is seen in the third panel to be localised around the VLT spiders,




Figure 3.6: Construction of simulated DTTS images with (lower panels) and without (upper panels) a 400 nm amplitude LWE. The first column shows the ideal stellar images with only the LWE, the second column includes XAO-corrected atmospheric residuals, the third column adds NCPES and the fourth column shows the final noisy DTTS images, re-scaled to simulated ADU counts.




Figure 3.7: Closed-loop performance of F&F for the LWE wavefront shown in Fig. 3.2 with simulated SPHERE/DTTS observing parameters described in Sec. 3.3.1. *Left:* 15s integrations of the final pre-coronagraphic DTTS PSFs with and without F&F correction (bottom and top respectively). *Right:* Strehl ratio of individual 0.1s DTTS exposures as a function of simulation time, showing significant improvement made by F&F (blue) over the un-corrected wavefront (red).




Figure 3.8: Wavefront phase maps illustrating the corrective ability of the first 50 Zernike modes, in a closed-loop simulation involving only LWE aberration. *Left:* Input LWE wavefront. *Middle:* Final corrective phase on deformable mirror which matches well the tip/tilt components of the wavefront. *Right:* Residual wavefront error after correction, seen to be dominated by high-frequency errors around the spiders, and additional high-order waffle and faulty actuator patterns present in the initial wavefront.

in addition to a weak waffle pattern and the influence functions of faulty actuators present in the initial wavefront.

3.4 Discussion

In these simulations F&F was found to have an effective closed-loop dynamic range capable of correcting an existing LWE aberration with discontinuity amplitude of 600 nm across the spiders, corresponding to a wavefront phase variance of $\sigma_\phi^2 = 1$ radian. This is twice the typical weak-phase approximation criterion $\sigma_\phi^2 \leq 0.5$ radians (Polo et al. 2013; Gonsalves 2001) due mostly to the inclusion of the second-order Strehl correction factor in Eq. 3.2. This increased dynamic range is in itself sufficient to cope with the majority of the fully-evolved cases of the LWE seen on-sky at SPHERE. Any further convergence issues may be simply avoided by ensuring F&F correction is activated as soon the LWE passes a critical detection threshold, in the same manner as Fig. 3.7. Alternatively, it is hoped that F&F will prove sufficiently stable to run continuously in the nominal SPHERE observing mode without manual intervention, such that it is possible to prevent the LWE from arising in the first place.

Spatial filtering of the wavefront estimates produced by F&F is necessary in order to remove high-frequency noise and artefacts of pupil asymmetry generated by the algorithm, in addition to providing an approximation of the corrective order of the chosen deformable mirror. The majority of the algorithm works analytically at the pixel level without any assumption of wavefront geometry or spatial order; it is therefore prone to propagating pixel-to-pixel noise in the input DTTS images through as high-frequency noise the final wavefront estimates. The re-formulating of Eq. 3.5 to include ϵ following the approach of Gonsalves (2001) provides a level of noise reduction for the odd wavefront, however the solution of v is significantly more prone to pixel-to-pixel noise due to the independent estimation of signs of the focal-plane field using Eq. 3.7. Filtering has been found to be effective at removing these noise sources, although other approaches such as amplitude apodisation of the input images has also been seen to reduce the severity of noise propagation (Korkiakoski et al. 2014). F&F operates under the implicit assumption of a completely even aperture function

A, hence any asymmetric structures such as VLT spiders result in high-frequency, low-amplitude phase ‘ghosts’ located at the conjugate positions of the amplitude masking element in the raw F&F output wavefront estimates. While such effects were not found to grow significantly in amplitude in closed-loop in cases where no spatial filtering techniques were used, they nonetheless degrade wavefront quality by a small amount and are best removed from final estimates.

It was also an initial concern that SPHERE’s main 41×41 actuator deformable mirror would be unable to cope with the sharp wavefront discontinuities present in all LWE wavefronts, hence the use of a 50-mode low-order basis throughout this paper to remove the high-frequency components of F&F estimates. As presented by Sauvage et al. 2016 (Sauvage et al. 2016) however, this may in reality be highly conservative as SAXO is expected be able to accurately implement a custom projection basis describing independent piston, tip and tilt errors across each VLT pupil segment and, if desired, also include a secondary smooth basis such as Zernike modes for independent control of NCPEs. A limited correction basis tailored to the LWE is also expected to help maintain algorithm stability by minimising the number of degrees of freedom available for the secondary loop to control. It is still worth noting however that the majority of the LWE may be eliminated by controlling only the first 50 Zernike modes as shown in Fig. 3.8, should a solution implementing only low spatial frequencies prove to be necessary for practical implementation in SPHERE.

Additional outstanding practical considerations not addressed by these simulations include:

- 1) **XAO loop conflicts:** The potential for conflict of F&F commands with the existing adaptive optics operations loop may be avoided by applying these corrections as reference offsets to the Shack-Hartmann sensor in the same manner as is proposed for NCPEs in Sauvage et al. 2011 (Sauvage et al. 2011). Further investigation is required to verify that the additional stroke used in compensating the LWE does not significantly degrade the performance of the main atmospheric correction loop.
- 2) **DTTS windowing modes:** Although the Hawaii I detector of the DTTS consists of an array of 1024×1024 pixels, it is usually run in a highly windowed mode to increase readout speed which cuts off the majority of the high-order PSF structure. However, it is expected that the F&F algorithm is able to cope with this reduction of field of view through appropriate normalisation, given that the majority of information beyond $\sim 3\lambda/D$ is dominated by pixel-to-pixel noise in the simulations presented in this paper.
- 3) **Dead Actuators:** SPHERE currently has a number of known dead or faulty actuators, which modify the phase response of the mirror and are typically masked out in amplitude during science operation. In order to provide the correct phase diversity to F&F, it is important to ensure an up-to-date DM model is used for estimating Φ_d .

3.5 Conclusions

The ‘Fast and Furious’ (F&F) algorithm is a promising solution for the real-time detection and correction of the low-wind effect (LWE), which is also in principle capa-

ble of compensating quasi-static non-common-path errors in part of the near-infrared optical path for the IRDIS and IFS instruments. This is achieved using no additional hardware and functions entirely within the nominal observation mode of the instrument by using non-coronagraphic images supplied by the existing DTTS imager present within the NIR instrument arm. Realistic closed-loop simulations show that F&F can robustly eliminate the characteristic ‘Mickey Mouse ears’ associated with a LWE-affected PSF within five iterations at 10 Hz cadence, providing significant recovery of Strehl and hence improvement in coronagraph performance under the best observing conditions. Future work will focus on validating F&F on a representative experimental optical testbench and, if successful, its implementation within the SAXO extreme adaptive optics system of SPHERE. Such algorithms may also prove useful in the long term for overcoming similar (or as-yet unforeseen) issues with XAO systems on ELT-class telescopes, where additional complications from mirror segmentation and associated phasing errors may lead to similar effects.

