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

Author: Wilby, M.J.
Title: Painting with starlight : optical techniques for the high-contrast imaging of exoplanets
Issue Date: 2018-11-27

Chapter 5

These are not the voids you're looking for

Using Ceres to calibrate SPHERE-IRDIS coronagraphic DPI observations of TW Hydrae

Michael J. Wilby, Jos de Boer, Rob G. van Holstein, Christian Ginski, Julien Girard, Anthony Boccaletti, and Christoph U. Keller

In preparation

Abstract

The use of coronagraphic optics and dual-band polarimetric imaging (DPI) techniques to remove unwanted stellar contamination from high-contrast images leads to significant transmission losses at small angular separations. This makes it challenging to determine whether features observed in these regions are due to intrinsic disk features such as inner cavities, or are simply artefacts of the imaging system. This work proposes two calibration approaches of differing levels of complexity, which can be used to correct the relative photometry of circumstellar disk observations. The first uses measurements of the radial coronagraphic extinction profile to directly normalise observations, while the second uses a full model of the imaging system to account for non-linear effects when forward-modelling the object in question. e analyse observations of the minor planet Ceres in order to measure the extinction profile of the SPHERE-IRDIS apodised Lyot coronagraph. From this an effective model of the coronagraphic imaging system is developed, which is used to calibrate polarimetric observations of the TW Hydrae protoplanetary disk. We find that a four-component power law model provides an accurate reconstruction of the observed TW Hydrae signal below 600 mas, and confirms that the innermost gap feature below 110 mas observed by van Boekel et al. (2017) cannot result from instrumental effects alone. The results of both calibration approaches are found to be consistent down to 75 mas, below which point the normalisation calibration is dominated by diffracted instrumental polarisation artefacts. Observing solar system bodies such as Ceres provides a useful approach for calibrating coronagraphic imaging systems, especially when combined with optical instrument modelling. We also find that for DPI observations, signal losses due to convolutional depolarisation effects are of equal importance to coronagraphic extinction, especially between angular separations of 100-300 mas. This factor should therefore always be accounted for before the detection of an inner disk feature is claimed.

5.1 Introduction

High-contrast, high-resolution scattered light imaging at near-infrared (NIR) wavelengths has in recent years developed into an incredibly versatile tool for observing protoplanetary disks, providing constraints on the disk morphology by allowing us to distinguish spiral arms (e.g., Benisty et al. 2015) and rings and gaps that yield the surface-height profile (de Boer et al. 2016; Ginski et al. 2016) and micron-scale dust composition (Muro-Arena et al. 2018) of the scattering surface of these objects. With ALMA now providing comparable spatial resolution in the sub-mm, it is now also possible to use multi-wavelength studies to place strong constraints on the vertical structure and decoupling between small and large grains due to pressure traps within of these disks (Pinilla et al. 2015), and hence build a more complete picture of the planet-forming environment (e.g. Marino et al. 2015; de Juan Ovelar et al. 2016). Dual-band polarimetric imaging (DPI, see e.g. Kuhn et al. 2001) has in particular become an extremely popular observing mode for scattered-light disk science. This is due to its capacity for effective starlight rejection, whilst also introducing minimal image artefacts when compared with the angular differential imaging (ADI, Marois et al. 2006) family of reduction algorithms, which tend to erase or distort extended disk signal.

Transitional disks are of particular interest for understanding the processes of disk evolution and planet formation. These objects are classified based on the morphology of their spectral energy distribution (SED), which displays a characteristic suppression of the disk's infra-red excess at NIR wavelengths, indicating the absence of the hottest dust components (e.g. Strom et al. 1989; Calvet et al. 2002; Merín et al. 2010). This is interpreted as being caused by dust depletions, or cavities, in the inner regions where disk material is being cleared in an inside-out fashion by radiation pressure from the central star (Clarke et al. 2001; Alexander & Armitage 2007). This interpretation has been confirmed for multiple transition disks (e.g. HD 142527 by Canovas et al. 2013, HD 169142 by Quanz et al. 2013b, and HD 135344B by Garufi et al. 2013).

Observations of such disks have revealed a wide variety of morphological features including multiple rings and gaps (e.g. RX J1615: de Boer et al. 2016, HD 97048: Doucet et al. 2007; Ginski et al. 2016, and HD 141569: Biller et al. 2015; Perrot et al. 2016); spiral structures (e.g. MWC 758: Grady et al. 2013; Benisty et al. 2015, HD 100453: Wagner et al. 2015; Benisty et al. 2017 and HD 135344B: Muto et al. 2012; Garufi et al. 2013); and potential exoplanet accretion sites (e.g. HD 100546: Quanz et al. 2013a, 2015; Currie et al. 2015, LkCa 15: Kraus & Ireland 2012; Sallum et al. 2015, and most recently PDS 70: Keppler et al. 2018). It is important to understand whether such features are indeed indicators of ongoing planet formation, or whether they can be explained by unrelated processes operating within the disk.

Since the existence of a central cavity is the hallmark of the transition process, there is a strong incentive to probe the innermost regions of protoplanetary disks in order to identify such features. However, a nearly all high-contrast imaging observations make use of attenuating coronagraphs to suppress on-axis starlight, which produce distinctive but often poorly characterised signal losses in the vicinity of the inner-working-angle (IWA[1]) of the instrument. Such instrumental features may easily be mistaken for a gap or cavity if not correctly accounted for, especially in low-inclination disks where the two can appear functionally identical. Even when

---

[1] For observations using a focal-plane coronagraph, this is commonly defined as the angular separation at which the instrument throughput drops to 50% of the wide-angle baseline.

analysing non-coronagraphic polarimetry data, the influence of an extended instrument point-spread function (PSF) can lead to significant, non-linear erasure of disk signal at angular separations interior to a few hundred milli-arcseconds. This "PSF smearing" effect was originally noted in Avenhaus et al. 2014a: here we refer to it as "convolutional depolarisation", and discuss the effect extensively in Sec. 5.4 onwards.

It is therefore desirable to develop a general calibration approach which can accurately recover the relative surface brightness of the disk as close to the IWA limit of the imaging system as possible. In this work we investigate the various instrumental factors which can result in artificial flux attenuation at small angular separations, with a view to calibrating observations made with the Spectro-Polarimetric High-contrast Exoplanet REsearch instrument (SPHERE, Beuzit et al. 2008). Particular attention is paid to the most popular coronagraphic and polarimetric imaging modes of the InfraRed Dual-beam Imager and Spectrograph (IRDIS) subsystem (Dohlen et al. 2008).

To this end we obtain H-band IRDIS coronagraphic observations of Ceres, in order to measure the two-dimensional extinction profile of the SPHERE-IRDIS Apodised-Pupil Lyot Coronagraph (ALC) for a bright, extended target of known flux distribution without the complication of stellar residuals (Sec. 5.2). The data is then used to validate a model of the coronagraphic imaging system, which allows analysis of the various contributions to measured flux as a function of angular separation (Sec. 5.3). A recipe for calibrating coronagraphic disk observations is then developed and validated using the Ceres data and disk models (Sec. 5.4), which includes a discussion of additional effects particular to DPI observations. This calibration procedure is then demonstrated for the case of TW Hydrae (Sec. 5.5), using the coronagraphic H-band DPI observations of van Boekel et al. (2017).

5.2 Ceres observations and data reduction

5.2.1 Observations

Calibration observations of Ceres were taken on the night of the 13th December 2016 as part of a wider effort to characterise the SPHERE NIR coronagraphic system. These observations were performed with IRDIS in H2/H3 dual-band imaging mode ($\lambda_c = 1588.8/1667.1$ nm, $\Delta\lambda = 53.1/55.6$ nm), and using the N_ALC_YJH_S apodised-pupil Lyot coronagraph (ALC) configuration. This mode is particularly useful for a first calibration attempt as it is currently the most commonly used SPHERE coronagraph for disk science observations.

At the time of observation, Ceres displayed an angular diameter of 0.6", a V-band magnitude of +8.3 and a phase angle of 17.6° (JPL horizons database, Giorgini et al. 2015). This makes it a perfect target for coronagraphic calibration as it provides an extended, near-uniform target without the complication of a bright central point source. With a diameter of 0.6", the disk of Ceres extends beyond the 0.2"-radius sphere of influence of the ALC coronagraph, whilst still being significantly smaller than the 1.0" angular resolution of the 40×40-lenslet SPHERE-SAXO Shack-Hartmann wavefront sensor (Sauvage et al. 2014), required for the XAO system to effectively use it as a natural guide star.

Sequences of 8 s exposures were taken in two positional configurations: firstly with the disk of Ceres centred on the coronagraphic axis, and secondly with the target offset from the ALC mask to provide a non-coronagraphic reference image. The on-axis dataset consisted of five exposure sequences totalling 9.6 minutes of integration time, during which the IRDIS detector was dithered in a 1×1 pixel square grid pattern. The off-axis configuration comprised two sequences totalling 2.1 minutes, at a pointing offset of 0.4”-0.5” with respect to the on-axis dataset. All data were taken in pupil-tracking mode in order to stabilise the coronagraphic optics, resulting in 3° of field rotation during the observing sequence.

Such a method of obtaining the coronagraphic extinction profile is complementary to another approach carried out in parallel, where a disk-less star (i.e. a simple point source) is observed and slowly shifted out from under the coronagraph mask by applying tip/tilt reference offsets via the AO system’s Shack-Hartmann wavefront sensor. The extinction profile is then reconstructed from the peak throughput of each image at the corresponding angular separation from the coronagraph axis. The results of this calibration can be found in the SPHERE instrument manual (ESO 2018) for IRDIS H-band and IFS H- to J-band measurements. The two methods are highly complementary, as this “point-source” method provides an estimate of the direct throughput of the coronagraphic system as a function of angular separation from the optical axis. On the other hand, the “extended-source” method using Ceres inherently includes non-linear diffraction and scattering contributions from all resolution elements of the target. The comparison of both approaches therefore highlights those regions where the extinction curve is no longer dominated by direct (useful) throughput for a particular extended target. The point-source method can also be sensitive to temporal variations in image Strehl ratio and errors in positioning accuracy, and provides only a one-dimensional estimate of the coronagraphic extinction profile. By providing a full image of the profile using only two images, the extended-source method is significantly more robust against these effects. Such an approach is also necessary to properly characterise asymmetric coronagraphic optics, such as the wedge coronagraphs of HST-STIS (Heap et al. 2000) and JWST-NIRCam (Krist et al. 2009), which cannot be calibrated using a simple 1-D radial profile.

5.2.2 Data reduction

All Ceres H2/H3-band images were first appropriately dark- and flat-corrected. Each frame was then independently corrected for sky background signal and frame-to-frame readout striping artefacts via column-wise median subtraction, using areas containing no significant Ceres signal. Sky frames were available but were not used in this case, as it was seen that they retained latent image signal from the previous object exposures. While negligible when compared to the total flux of Ceres (less than 0.2%), this latent signal was seen to create minor but noticeable artefacts beyond the edge of the Ceres disk when used in the image reduction process. A mask of the many hot and rogue pixels of the IRDIS detector was however created using the dark/flat-calibrated sky frame data, based on a manually optimised threshold chosen to capture all bad pixels whilst avoiding overly masking additional pixels.

On-axis frames were then shifted to correct for the 1×1 px dither pattern used during the observation sequence, before extracting the H2 and H3 channels from the image. The stability of the ALC focal-plane mask with respect to the IRDIS detector plane is extremely high, making it unnecessary to perform any additional sub-pixel correction when co-aligning the coronagraphic extinction profile of each image. How-

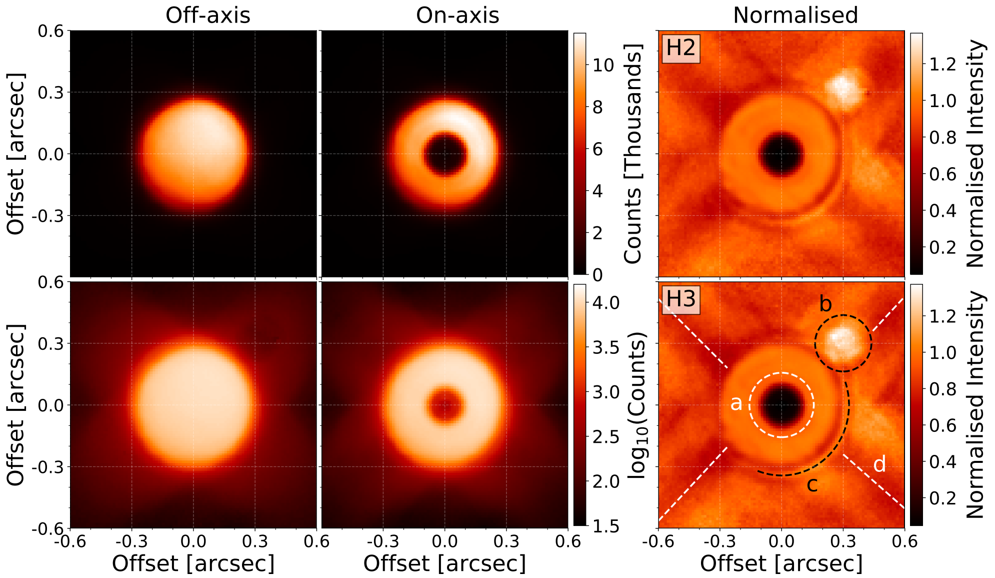

Figure 5.1: Reduced H2-band images of Ceres calibration observations, and the normalised ALC throughput profiles in both the H2 and H3 bands. *Left column:* non-coronagraphic images, with the ALC focal-plane mask located an average of 0.42” to the top-right of the Ceres disk edge. *Middle column:* coronagraphic observations with the focal-plane mask centered on the Ceres disk. The bottom row shows the log-scaled equivalent of linearly scaled top row, showing the inner regions of the coronagraphic mask and modifications to the diffraction pattern beyond the disk edge. *Right column:* linearly-scaled ratios of on-axis and off-axis intensity images, which correct for spatial variations in the Ceres surface brightness profile and are used to measure coronagraphic throughput curves of Fig. 5.2.

ever, due to relatively poor observing conditions there were significant frame-to-frame variations in image quality, which mostly manifested in changes in the shape of the first Airy ring and hence the centroid location of the target with respect to the ALC mask.

To account for these relative positional offsets, the central coordinates of each frame in both the on-axis and off-axis image sets was registered by fitting a model Ceres profile using the Levenberg-Marquardt least-squares fitting routine of the astropy python package. The observed Ceres surface brightness profile was found to be accurately modelled as a spherical Lambertian reflector with the appropriate illumination angle, convolved with a simulated un-occulted ALC instrument PSF. This model was used as the least-squares fitting function, with central coordinates and peak flux allowed to vary as free parameters. The central 175 mas of on-axis images was first masked out, in order to avoid biases due to the coronagraphic extinction profile.

The off-axis image frames were then co-aligned to sub-pixel accuracy and median-stacked, to form a master non-coronagraphic image reference for the Ceres disk. This profile is shown in the left panels of Fig. 5.1. The master off-axis reference image was then separately co-aligned with each on-axis image after correcting for the average field rotation angle of each exposure sequence, then used to individually normalise each on-axis frame. De-rotation of each frame within an exposure sequence, or when co-aligning the off-axis reference images, was found not to be necessary: the low amount of field rotation in this dataset meant that pixels at the edge of the Ceres disk shifted by less than 0.2 pixels during individual exposure sequences.

The final coronagraphic extinction profile image was then obtained via median-stacking of the normalised dataset: this profile is shown in the right-hand side of Fig. 5.1 for both the H2 and H3 bands. This frame-wise reduction approach provides the optimum normalisation of the edges of the Ceres disk under non-ideal observing conditions, thereby maximising the useful field over which the coronagraphic extinction profile can be analysed without the influence of disk-edge effects.

5.2.3 Analysis

5.2.3.1 Image features

The circular extinction profile of the coronagraph is clearly visible in the center of the Ceres normalised intensity images (feature *a* in the bottom-right panel of Fig. 5.1). Both H2 and H3-band images were seen to be nearly identical, save for some small increase in scale of the diffraction pattern beyond the edge of the Ceres disk, as would be expected for the small change in wavelength between bands. The disk of Ceres is seen to be well corrected for surface brightness variations out to an angular separation of approximately 250 mas, which allows for a straightforward extraction of the full ALC extinction profile.

Numerous secondary features are also visible in the outer regions of these normalised images, which can all be explained with reference to the un-normalised data. Firstly, feature *b* is an artefact of the ALC focal-plane mask blocking part of the spider diffraction structure in the off-axis images (visible in the bottom-left panel of Fig. 5.1). The asymmetric arc-like feature, labelled feature *c*, is attributed to a coma-like quasi-static aberration in the system which varied between the on-axis and off-axis datasets. By changing the morphology of the first Airy ring, this aberration correspondingly alters the brightness distribution of the outer regions of the Ceres disk in the same way as discussed in Sec. 5.2.2. In this case, the effect results in an over-normalisation in the top-left half of the image and a corresponding under-normalisation in the lower-right region. The four dark channels in the spider diffraction structure, labelled feature *d*, is a result of occultation of the central regions of the Ceres disk in the on-axis images, as is visible in the lower-middle panel of Fig. 5.1. This is because the majority of this diffraction structure is produced by the spider mask component of the Lyot stop, which sits downstream of the focal-plane mask and therefore diffracts less flux along the diagonals corresponding to the occulted regions of the disk.

5.2.3.2 Radial extinction profile

The radial coronagraph extinction profile for these observations was extracted after precisely determining the location of the optical axis in the image, since this is not known a priori. Due to the high level of azimuthal uniformity seen in the ALC extinction region normalised images, the centering was refined via an iterative grid search which minimised the variance of the extracted radial profile for angular separations below 200 mas. The resulting extinction curves are shown for both H2 and H3 bands as the solid coloured lines in Fig. 5.2. It can be seen that the coronagraph profile has extremely low variance in the regions of significant extinction (<125 mas), while the standard deviation of the outer regions of the profile which is not influenced by disk-edge effects (<250 mas) remains below the 2% level. We can therefore conclude that the extinction profile of this SPHERE-IRDIS ALC coronagraphic system is az-

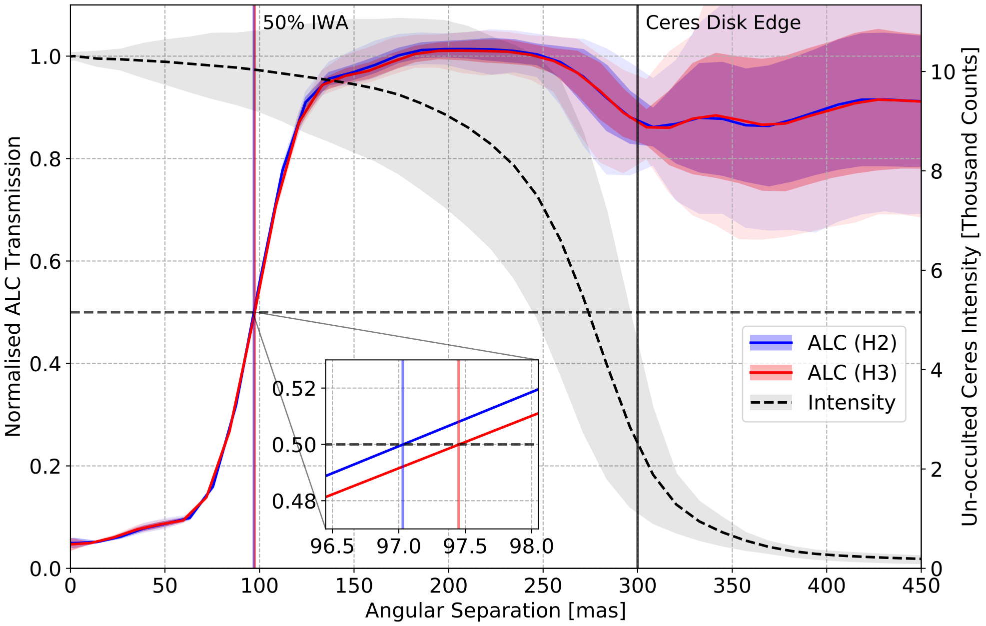

Figure 5.2: Radial ALC coronagraph extinction profile, as measured from Ceres data (right column of Fig. 5.1). Blue and red curves denote the median azimuthally averaged normalised intensity profile for the H2- and H3-bands respectively, with dark shading denoting the 1σ azimuthal variability and light shading denoting the extrema of the data. The grey dashed curve denotes the un-occulted Ceres surface brightness profile (left column of Fig. 5.1), with shading denoting the extrema. The edge of the Ceres disk is located at 300 mas, and the 50% IWA contour is located at approximately 97 mas, denoted by red and blue vertical lines in the inset panel for the H2- and H3-bands, respectively.

imuthally uniform to a high degree, facilitating an accurate calibration using a 1-D radial profile in this case.

The 50% IWA is measured to be 97.0 ± 0.3 mas and 97.5 ± 0.3 mas for H2 and H3 respectively, which is within 2 mas of the value reported by the results of the point-source offset method (ESO 2018). Note that this value is higher than the physical 92.5 mas radius of the coronagraph mask which is often incorrectly taken as the masking radius for analysis of coronagraphic data: this figure clearly shows that the coronagraph is responsible for significant extinction out to approximately 200 mas, which if not correctly normalised may be erroneously claimed as the outermost edge of a central cavity in an equivalent protoplanetary disk observation.

It can be seen that in Fig. 5.2 the profile levels off in the inner regions behind the coronagraph mask (<75 mas), tending to a value much greater than the $< 5 \times 10^{-3}$ on-axis H-band extinction ratio which was measured for this coronagraph during commissioning laboratory testing (Guerri et al. 2011) and on-sky (Fusco et al. 2016). This is behaviour is not seen in direct throughput curve measurements made via the point-source calibration approach, since the additional signal is produced by diffraction from all resolution elements of the target outside the ALC mask. Accurately reproducing this non-linear contribution to the total signal for an extended target is one of the major motivators of the coronagraphic system modelling presented Sec. 5.3, since it ultimately determines the smallest angular separation at which useful astrophysical signal can be recovered.

Additional modelling is also required to determine why the region between 175-

250 mas exceeds a normalised intensity of unity: this should not be possible if the coronagraphic mask is the only source of extinction. It was found however that this can be effectively explained by variations in Strehl ratio between the on- and off-axis images: this consideration is presented in detail in Sec. 5.3.3 with the aid of a coronagraphic system model.

5.3 Modelling the SPHERE ALC Coronagraph System

5.3.1 Description of the coronagraphic system model

It is clear from Sec. 5.2.3 that a model of the ALC system is required in order to fully explain the observed extinction curve, and to facilitate an optimal correction scheme for science images. For this purpose we use a three-plane Fourier propagation model which incorporates the amplitude apodiser, focal-plane mask and Lyot stop designs shown in Fig. 5.3. These planes are based on the comprehensive description of the ALC design and laboratory testing provided by Guerri et al. (2011). The additional structures visible in the Lyot stop pupil plane were added during commissioning of SPHERE, and mask dead actuators on the SAXO deformable mirror (Dohlen et al. 2016).

A coronagraphic image $I(x, y)$ is produced by convolution of the surface brightness distribution of the astrophysical object $O(x, y)$ with the spatially variable PSF $p_{ij}(x, y)$ which describes the attenuation of the coronagraphic system. This can be written in the form of a discrete convolution, as

$$I(x, y) = \sum_{ij} O(i, j) p_{ij}(x, y), \tag{5.1}$$

where the indices i and j refer to the discrete sampling points in the focal plane along the x and y axes respectively, which functionally correspond to pixels in the detector array. In this work, pupil plane coordinates are denoted in (u, v) space while focal-plane coordinates are denoted in (x, y) space.

The individual point-spread functions required for this computation are calculated by re-imaging the respective pupil and focal-planes of the coronagraph, such that

$$p_{ij}(x, y) = \left| \mathscr{F}\left[L(u, v)\ \mathscr{F}\left[m(x, y)\ \mathscr{F}\left[A(u, v) \mathrm{e}^{i\phi_{ij}(u,v)} \right]\right]\right]\right|^2 . \tag{5.2}$$

In this equation $A(u, v)$ denotes the pupil amplitude apodisation function (panel 1 of Fig. 5.3), $m(x, y)$ the focal-plane mask (panel 2), and $L(u, v)$ the downstream Lyot stop (panel 3). The location of the PSF in the focal plane is controlled by adding a tip-tilt wavefront error $\phi_{ij}(u, v)$ to the first pupil plane, which is calculated to exactly match the location of the desired sampling point ij. The Fourier transform operation is denoted by $\mathscr{F}\left[\right]$, implemented here using the array-based Fourier transform routine described in Soummer et al. (2007b) which is suited to the efficient computation of Lyot coronagraphic PSFs.

The intermediate focal-plane containing the coronagraphic mask is sampled at a fixed resolution of 5 pixels per λ/D, as a trade-off between speed and resolution when computing a large grid of PSFs, each with a significant final field-of-view. The

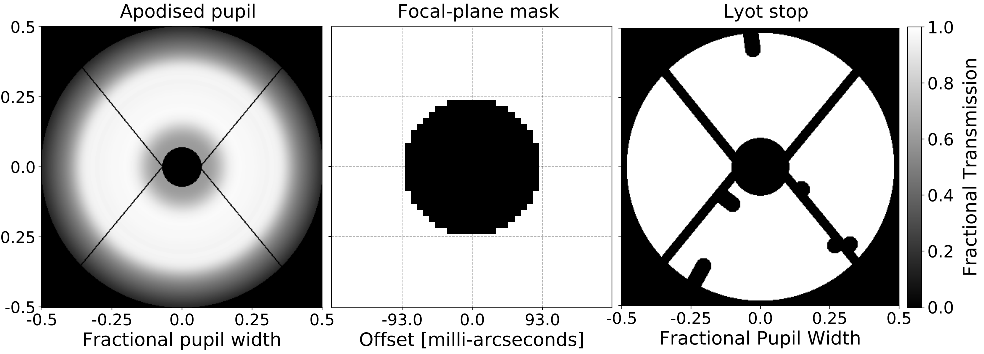

Figure 5.3: Three-plane ALC model used for optical system modelling, matching the SPHERE-IRDIS N_ALC_YJH_S ALC coronagraph configuration (ESO 2018). The VLT pupil geometry and ALC_2 Lyot stop design, including dead actuator masks, is based on Fusco et al. (2014). The APO1 pupil apodisation function is based on the H-band laboratory measurements of Guerri et al. (2011). The focal-plane mask is computed at a resolution of 5 pixels per λ/D.

simulation wavelength can be freely adjusted by varying the effective radius of the focal-plane mask and appropriately scaling the sampling grid of the final image plane. Phase aberrations from sources such as XAO-filtered atmospheric residuals and non-common path errors can also be added to either the upstream or downstream pupil planes, such that $A(u, v)$ and $L(u, v)$ may become complex in Eq. 5.2.

5.3.2 Diffraction-limited model

In order to directly compare with observations, we evaluate our ALC model with an object model $O(x, y)$ emulating the Ceres surface brightness profile at the time of observation. This consists of the spherical, perfectly Lambertian reflector model that was previously used for image registration in Sec. 5.2.2, with a phase angle of 17.6°. A datacube of diffraction-limited ALC imaging PSFs $p_{ij}(x, y)$ was then computed according to Eq. 5.2 for each resolution element of the IRDIS focal plane where $O(x, y)$ is non-zero, at the H2-band central wavelength of 1588.8 nm. An equivalent un-occulted datacube was also computed by omitting the focal-plane mask and directly computing the Fourier transform squared of $L(u, v)A(u, v)e^{i\phi_{ij}(u,v)}$, to act as a reference for normalisation. The Ceres observations presented in Sec. 5.2 were then emulated by computing the spatially-variable convolution of Eq. 5.1 with both the occulted and un-occulted $p_{ij}(x, y)$ datacubes, and normalising the first resulting image with the second.

Figure 5.4 shows the normalised intensity radial profile of this simulated image model (blue curve), which is directly comparable to the H2-band Ceres observations of Fig. 5.2, shown here as black data points with grey shading. Here the Ceres curve has been normalised to unity between 200-250 mas to properly account for the effects of variable seeing conditions during the observation window: this approach is justified in Sec. 5.3.3 where such factors are explored in detail.

The top panel of Fig. 5.4 shows that this diffraction-limited ALC model provides an excellent fit to the measured Ceres extinction profile, which is accurate to better than 2% for all angular separations above the 50% IWA and better than 5% up to the 92.5 mas focal-plane mask edge. The IWA itself is predicted to be 98.1 ± 0.6 mas, only 1 mas larger than that measured on-sky. This close agreement is especially impressive given that we use here the simplest possible forward model of the optical system,

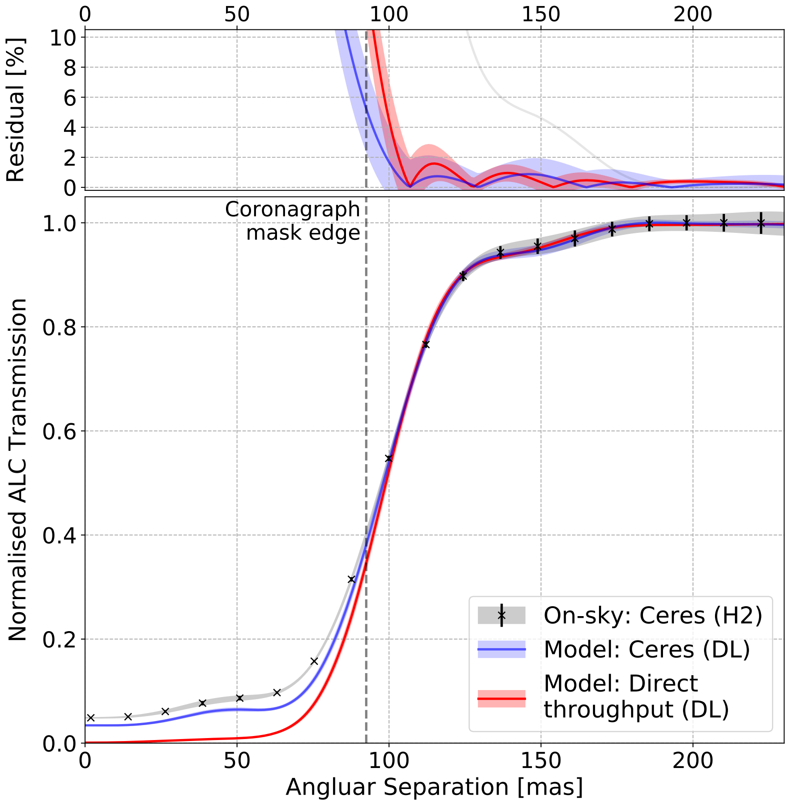

Figure 5.4: Comparison of the diffraction-limited (DL) ALC imaging simulation with the on-sky Ceres H2-band extinction profile measurement of Fig. 5.2. ***Bottom:*** Radial extinction profiles for on-sky Ceres data (black), simulated image of a Ceres-like object model (blue), and simulated direct throughput curve of the ALC system model (red) as a function of angular separation from the optical axis. ***Top:*** Absolute percentage residuals between both models and the Ceres on-sky data. Grey shading denotes the region in which systematic residuals are worse than the null hypothesis, where coronagraphic extinction effects are ignored.

which is fully constrained by the literature description of the IRDIS ALC coronagraph and hence has not been fitted to the Ceres observations at any stage. However, this model cannot fully account for all flux seen in the inner regions of the measured profile, which leads to an rapid increase in fractional error inside the IWA due to the small transmission values involved. This discrepancy could be due to a variety of instrumental effects such as non-common-path errors (NCPEs), and is addressed in Sec. 5.3.4.

Also plotted in Fig. 5.4 is the direct throughput of the coronagraphic system (red curve), measured as the relative transmission of pixel i, j for each coronagraphic PSF $p_{ij}(x, y)$ by analogy with the point-source calibration method discussed in Sec. 5.2.1. This direct throughput curve can be seen to tend towards zero at the coronagraphic axis, and is systematically lower than the full Ceres image model (blue curve) below approximately 110 mas. The difference between these two curves is simply the integrated signal that is diffracted inwards by the Lyot stop from un-occulted regions of the disk. It is noteworthy that there is still significant direct throughput for separations interior to the 92.5 mas edge of the focal-plane ALC mask. This would imply that it is in principle possible to extract useful flux information down to angular separations of approximately 60 mas for this coronagraph, however in reality this will almost always prove impossible due to the dominance of speckle noise or instrumental polarisation artefacts at these small angular separations.

It is this direct throughput metric which constituents the useful throughput of the coronagraphic system, since it is the fraction of signal that the optical system correctly maps from object $O(i, j)$ to image $I(i, j)$, and hence provides spatially-resolved information about the target. The most straightforward calibration approach is therefore to normalise the central regions of all disk observations using the radial direct-throughput curve of the ALC coronagraph presented here, to correct for the relative useful throughput at all angular separations. As can be seen from Fig. 5.4 however, such an approach will result in a systematic flux overestimate by a factor corresponding to the fractional diffracted-light contribution at each pixel. This leads to a substantial error at small angular separations, which for the case of Ceres is equal to the red curve in the upper panel of Fig. 5.4. This issue is further complicated by the fact that the relative contribution of non-useful diffracted signal is inherently source-dependent, based on the surface brightness profile of the object under observation. This makes it impossible to determine a single limiting angular separation above which the accuracy of the calibrated signal can be trusted for all observations using the same coronagraph. A proper calibration for these innermost regions is therefore only possible by using an accurate forward model of the coronagraphic imaging system, which can determine this on a target-by-target basis. To this end we extend the diffraction-limited model in the following subsections to consider the impact of various instrumental factors on the extinction profile, in order to try to minimise the discrepancy between our model and the Ceres observations.

5.3.3 Effects of residual atmospheric turbulence

The most obvious factor affecting on-sky images is the presence of residual atmospheric turbulence after correction by the SPHERE XAO system (SAXO, Fusco et al. 2016). The reason for modelling this is two-fold: to assess the impact of residual phase errors on the inner regions of the extinction profile, and to explain the morphology of the outer regions of the measured Ceres curve as noted in 5.2.3.2.

For this work, residual XAO-corrected phase-screens were generated using SPHERE-like AO simulation code originally developed in Wilby et al. (2016b), based on a Kolmogorov spatial power spectrum which is suppressed for frequencies controllable by the SAXO 41×41 actuator deformable mirror. The final functional form of this XAO-corrected spatial power spectrum Φ_{AO} was chosen to be

$$\Phi_{\mathrm{AO}}(k, r_0) = \begin{cases} 5\times10^{-3}\Phi_{\mathrm{Kol}}(k, r_0) + 0.5\Phi_{\mathrm{Kol}}(k_{\mathrm{crit}}, r_0) & k < k_{\mathrm{crit}} \\ \Phi_{\mathrm{Kol}}(k, r_0) & k \geqslant k_{\mathrm{crit}} \end{cases} \tag{5.3}$$

where Φ_{Kol} is the Kolmogorov phase power spectrum as a function of spatial frequency k (in units of cycles per pupil, or equivalently λ/D) and Fried parameter r_0. The spatial frequency cutoff k_{crit} is set to be $20\ \lambda/D$, matching the control radius of the SAXO deformable mirror: beyond this cutoff the power spectrum is identical to Φ_{Kol}. The resulting AO-corrected power spectrum is shown in the top-left panel of Fig. 5.5 (solid line), compared to the equivalent Kolmogorov power spectrum Φ_{Kol} (dashed line). Variable seeing conditions can be emulated using this model by simply increasing or decreas-

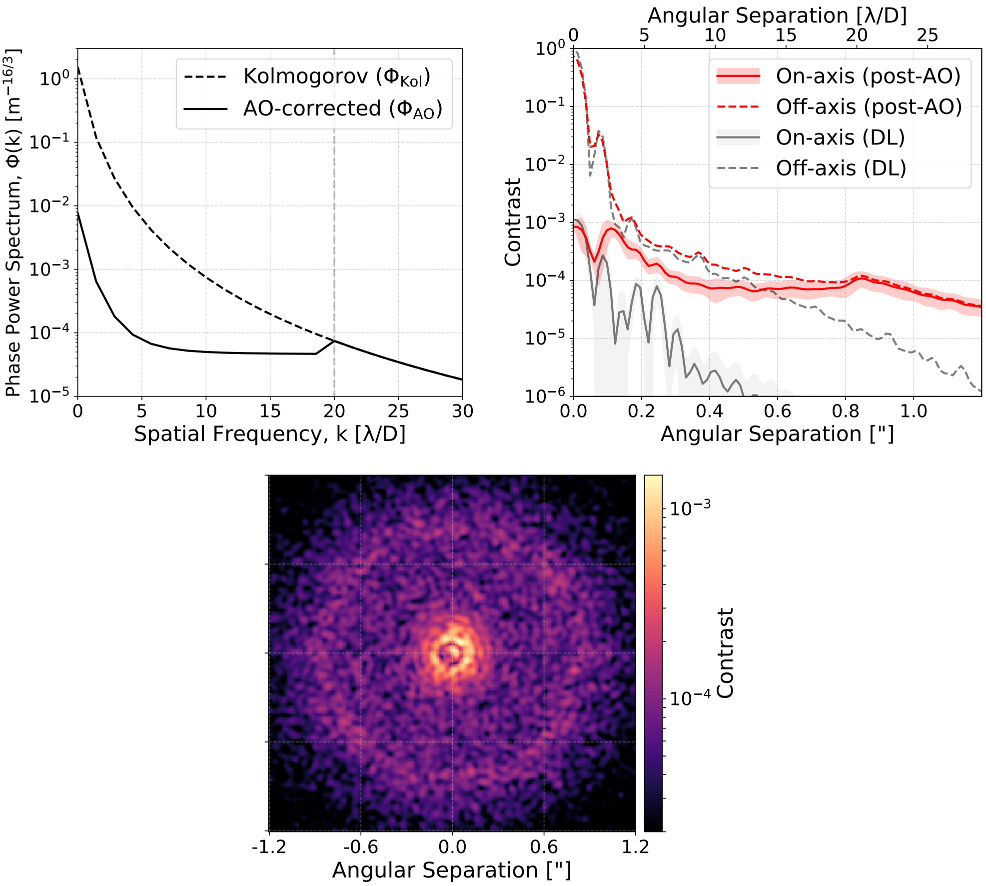

Figure 5.5: Key modelling parameters for SPHERE-like residual atmospheric phase errors. ***Top left:*** Power spectral density of XAO-corrected residual phase aberrations corresponding to Eq. 5.3 (solid line), with respect to purely Kolmogorov turbulence (dashed). ***Top right:*** The resulting raw coronagraphic contrast curve (on-axis, solid red line) and un-occulted PSF (off-axis, dashed red line), compared to the diffraction-limited case (DL, grey).***Bottom:*** Corresponding on-axis coronagraphic PSF, comprising the sum of ten individual speckle realisations.

ing r_0, and hence the overall image Strehl ratio, whilst keeping the AO spatial filtering function constant.

The functional form of this AO-corrected atmospheric spatial power spectrum was chosen to allow for independent modification of the fractional suppression amplitude of the Kolmogorov spectrum, and a contrast saturation threshold which is reached at high spatial frequencies. These parameters are tailored in addition to r_0 in order to reproduce the typical on-sky ALC contrast profile under median seeing conditions, as reported by Fusco et al. (2016). Replication of the exact system response of SAXO is not crucial to this investigation, which simply seeks to probe how SPHERE-like post-XAO atmospheric phase residuals qualitatively affect the coronagraphic extinction curve measurements of Sec. 5.2: this approximate model is sufficiently realistic for this purpose.

An example short-exposure PSF comprising an average of ten speckle-field realisations is shown in the bottom panel, in which both the 20 λ/D control radius and elevated low-frequency residual speckles can be seen, matching the main features of

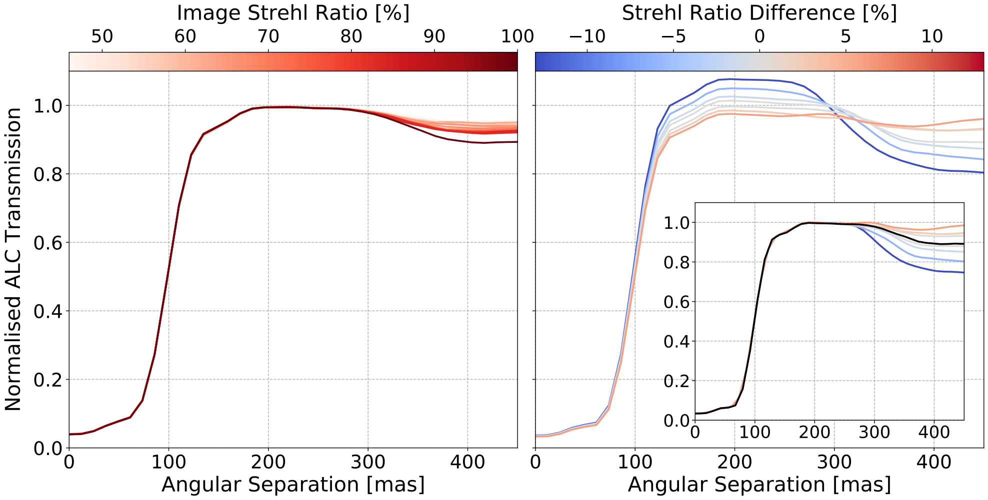

Figure 5.6: Simulation of the effect of varying Strehl ratio on the retrieved normalised intensity profile from Ceres observations. ***Left:*** Both on-axis and off-axis images are convolved with the same p_{ij}, and hence identical r_0 and seeing conditions. The darkest line corresponds to the diffraction-limited case. ***Right:*** On-axis image has a fixed Strehl ratio of approximately 75%, while the Strehl ratio of the off-axis image allowed to vary. A positive Strehl ratio difference corresponds to the off-axis image having a higher Strehl ratio. The inset panel shows the same data, normalised to unity between 200-250 mas. The diffraction-limited result is over-plotted in black, for comparison.

the chosen spatial power spectrum. The corresponding azimuthally-averaged contrast curve is shown in the top-right panel of Fig. 5.5 (solid red line), with the upper limit of the red shading denoting the 1σ sensitivity limit. Here the red dashed line shows the un-occulted (off-axis) PSF, while the grey curves show the equivalent diffraction-limited instrument performance for comparison purposes.

Figure 5.6 summarises the impact of these residual atmospheric phase aberrations on the coronagraphic extinction profile: the left-hand panel shows the simulated extinction profiles for various values of r_0, obtained in the same manner as in Sec. 5.3.2, where both the on-axis and off-axis images are produced under identical seeing conditions. This clearly shows that the inner coronagraphic extinction profile is not inherently dependent on residual atmospheric phase errors, with the only influence being minor variations beyond the 300 mas edge of the Ceres disk. This result can be anticipated, since the extinction profile in the science image is produced by near-diffraction-limited re-imaging of the coronagraphic focal plane mask itself. The morphology of this extinction profile is therefore dominated by diffraction from the oversized Lyot stop, making it insensitive to all phase errors upstream of the coronagraphic mask.

As the right-hand panel of Fig. 5.6 shows however, noticeable measurement artefacts are introduced when attempting to characterise the coronagraphic profile if even small variations in image Strehl ratio occur between the acquisition of the on-axis and off-axis datasets. These curves were computed by fixing the Strehl ratio of the on-axis image at 75%, and computing off-axis images for a variety of different r_0 values and hence image Strehl ratios. An average of ten instantaneous phase-screen realisations was used to produce each image, a number which was chosen because it provides some speckle averaging necessary to obtain a reliable estimate of image Strehl ratio for a given r_0, while still allowing relatively fast computation of the full spatial grid of aberrated PSFs. It was not necessary for this investigation to perform the extremely computationally expensive task of simulating all $p_{ij}(x, y, t)$ required in order to properly model the speckle statistics for a single $8\ s$ Ceres exposure: for a speckle coherence time of 1 ms this would require the computation of approximately 10^7 individual PSFs for each simulated Ceres image.

The best qualitative match to the measured Ceres normalised intensity curves of Fig. 5.2 is obtained for an off-axis image Strehl ratio only 1% lower than that of the on-axis image. Such a small change can easily be expected to have occurred during the one-hour Ceres observing window due to natural changes in seeing conditions. This differential image Strehl ratio effect naturally results in a profile that exceeds unity by a few percent between 200 to 250 mas, before dropping outside the disk edge due to over-normalisation by additional scattered light in the off-axis image at these separations. The quantitative agreement between these seeing simulations and the Ceres observations is however mediocre: this is to be expected given the simplified AO model used in this work coupled with the lack of inclusion of additional factors such as NCPEs, which as noted in Sec. 5.2.3.1 are most likely responsible for specific features seen at the edge of the Ceres disk.

Most importantly for this work however, there is a simple solution for correcting this imaging artefact: the inset panel in the right-hand side of Fig. 5.6 shows the same data as the main panel, simply normalised to unity by the flattest region of the profile between 200-250 mas. This normalisation correction neatly recovers the diffraction-limited profile below angular separations of 250 mas, independent of all seeing artefacts. The same correction is therefore applied to the Ceres extinction profile measurements from Fig. 5.4 onwards, in order to accurately estimate the true coronagraphic extinction profile and fairly compare with simulations.

5.3.4 Additional considerations

In order to properly account for the discrepancy between Ceres observations and the diffraction-limited imaging model shown in Fig. 5.4, it is necessary to include one or more additional effects internal to the instrument which contribute additional signal underneath the coronagraph mask. Physical effects which potentially fulfil this requirement include the presence of non-common path aberrations downstream of the focal-plane mask, a mis-alignment of the ALC focal-plane mask itself along the optical axis, or the high-frequency vibration of downstream optics: each of these cases are explored in detail in the following subsections. Low levels of instrumental scattered light can also produce noticeable signal where ALC extinction is greatest: this last effect is however difficult to model as it is inherently dependent on the full description of reflective surfaces in the instrument, making its treatment beyond the scope of this paper.

Since it is difficult to comment on the various scenarios without additional information, we perform numerous simulations to estimate the magnitude of each effect which is required to address the observed discrepancy. This is then compared to SPHERE performance budgets, and hence assess the likelihood of each. Fig. 5.7 shows the best fit models to Ceres observations for each of the three considered effects, based on minimised least-squares residuals for angular separations above 60 mas. All can be seen to make significant improvements over the diffraction-limited case.

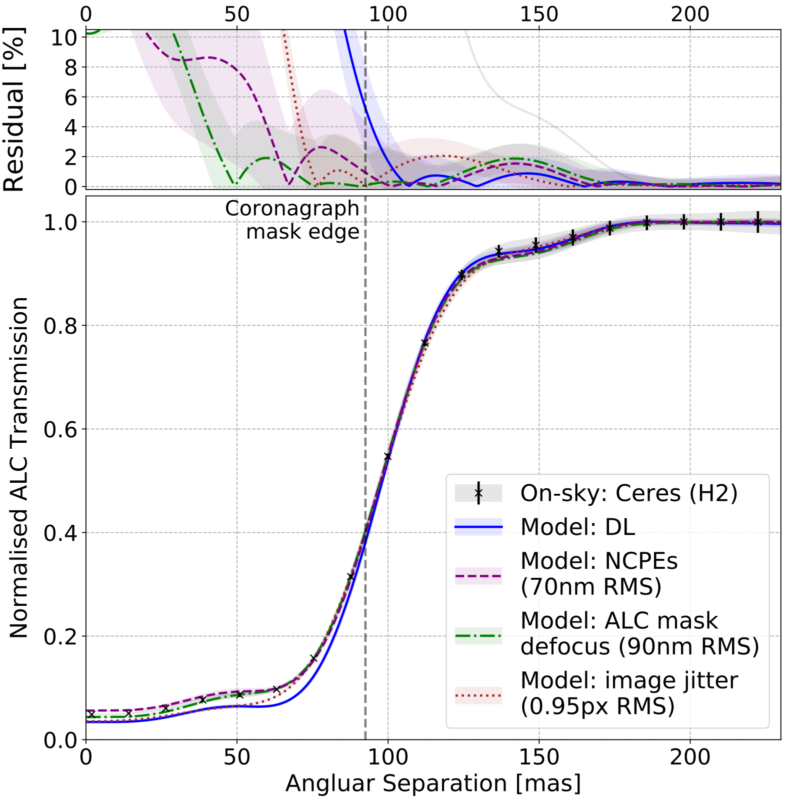

Figure 5.7: Comparison of the various non-diffraction-limited models presented in Sec. 5.3.4, compared to the on-sky Ceres extinction profile measurement. *Main panel:* ALC extinction profiles for H2-band Ceres measurements (black data points) and diffraction-limited image model (DL, solid blue curve) of Fig. 5.4. Also plotted are the best-fit image models which include the effects of non-common path errors (NCPE, dashed purple curve), ALC focal-plane mask defocus (dot-dashed green curve), and high-speed vibrational image jitter on the detector (dotted brown curve). *Top panel:* Percentage residuals between the aforementioned models and the measured Ceres extinction profile. Grey shading denotes the region in which systematic residuals are worse than the null hypothesis, where coronagraphic extinction effects are ignored.

5.3.4.1 Non-common path errors

Non-common path errors are present to some extent in all high-contrast imaging data, and contribute additional speckle noise which is strongest at small angular separations. If these aberrations occur downstream of the ALC focal-plane mask, re-imaging through the Lyot plane is no longer diffraction-limited and will correspondingly result in less efficient coronagraphic extinction. For SPHERE, these aberrations are typically estimated to contribute an average wavefront error of approximately 40 nm RMS (Fusco et al. 2016), which is consistent with the baseline requirement of 43 nm RMS for high-performance H-band operation. In addition, offline phase-diversity compensation routines have been measured to further reduce these NCPEs to 20 nm RMS using the internal calibration source (Sauvage et al. 2014). This performance budget does not however include the low-wind effect (Sauvage et al. 2016) which, while not itself due to NCPEs, significantly elevates image speckle noise under certain observing conditions.

We model these NPCEs by drawing randomised coefficients for the first 200 Zernike modes, with the resulting wavefront phase maps then spatially filtered in the Fourier domain to have a $1/f^2$-decreasing spatial power spectrum, as is often used to model NCPEs (Sauvage et al. 2007; Lamb et al. 2016). These randomised low-order phase aberrations are then injected in the downstream Lyot stop pupil plane, with five individual realisations being combined for each wavefront RMS value in order to somewhat average out speckle noise.

The best fit to Ceres observations was found to be for 70 nm RMS of $1/f^2$ NCPE, which provides a fit quality of better than 2% down to the 92.5 mas ALC mask edge, and better than 5% down to 60 mas. This is therefore significantly more accurate than the diffraction-limited model, being able to accurately account for the observed Ceres signal at all angular separations where the throughput of the ALC system is non-negligible. It should be noted that by requiring 70 nm RMS of NCPE to adequately explain the discrepancy, this scenario exceeds the 40 nm RMS performance tolerance reported during SPHERE calibration and testing. Such a 30 nm RMS increase in wavefront error is however not inconceivable, especially given that the presence of relatively strong, NCPE-like low-order aberrations was already noted during the reduction of Ceres data in Sec. 5.2.

By looking closely at the main panel of Fig. 5.7, it can be seen that the very innermost region of the profile below 60 mas is also slightly over-estimated compared to the region from 60 to 100 mas: this indicates that a $1/f^2$ spatial power spectrum may be marginally steeper than the optimal model for SPHERE NCPEs in this case, if it is indeed the underlying source of the signal discrepancy. While more complex NCPE models are likely to provide a significantly improved fit to the observed Ceres profile, the inclusion of more than a single additional free parameter is not justified here given the limited calibration data provided by a single measurement of the extinction profile. Most importantly, there is currently no information regarding the temporal variability of the ALC extinction profile, which makes it impossible to constrain a useful time-dependent NCPE model at this stage.

5.3.4.2 Out-of-focus ALC focal-plane mask

If the ALC focal-plane mask is accidentally mis-aligned along the optical axis, it will be out-of-focus with respect to the coronagraphic focal plane. This will increase transmission for resolution elements near the optical axis, because the mask now occults a marginally defocussed PSF which has more power at larger angular separations, hence a larger fraction lies outside the mask edge and is transmitted. Equally, resolution elements more than a diffraction-width from the mask edge will see a negligible change in throughput, making this an effect which only modifies the innermost regions of the extinction profile. The SPHERE longitudinal alignment tolerance for the focal-plane mask is ±1 mm, which is equivalent to 23 nm RMS of defocus aberration in the coronagraphic focal-plane (Guerri et al. 2011).

A shift of the ALC mask along the optical axis is modelled here by applying the corresponding focus error in the upstream apodiser pupil plane, and then removing it in the downstream Lyot stop pupil plane by applying the opposite phase. As can be seen in Fig. 5.7 the best model for this scenario provides an excellent fit to the observed Ceres extinction curve, with an accuracy of better than 2% down to angular separations of 40 mas. However, the 90 nm RMS of focus error required to produce this excellent fit corresponds to a 4 mm physical translation of the focal-plane mask from the nominal position (computed according to the optical description of Guerri et al. 2011), which is four times larger than the quoted alignment tolerance of this optic. It would be necessary to check the alignment of the SPHERE NIR coronagraph wheel in order to determine whether this is in fact a realistic scenario: such a mask defocus could easily be overlooked during instrument calibration, as it has a negligible effect on the extinction of an on-axis point source. The quality of fit which this model provides using only a single free parameter, combined with the fact that it is a time-independent effect, would make it an important factor to include in the optical system model were it shown to be the underlying cause of the discrepancy.

5.3.4.3 Image positional jitter

If optics downstream of the ALC focal-plane mask undergo mechanical vibrations or deformations on timescales of less than or equal to the exposure time, the coronagraphic image will jitter with respect to the detector plane, blurring the extinction profile by a corresponding amount. This effect is modelled here by shifting the diffraction-limited Ceres coronagraphic extinction profile model, using Gaussian-distributed displacement amplitudes in random directions to 1/50th of a pixel precision. An average of 200 individual realisations are combined for each RMS value, which when considering a single 8 s Ceres exposure corresponds a characteristic jitter frequency of 25 Hz.

Figure 5.7 shows that the best-fit for this jitter scenario still provides some improvement over the diffraction-limited case, although it is not as successful as either of the preceding effects in accounting for the overall morphology of the extinction profile. In particular, this positional jitter model cannot explain the observed flux excess for angular separations below approximately 75 mas. It should be noted that the best-fit model requires there to be a 0.95 pixel RMS jitter amplitude between the coronagraphic focal-plane mask and the detector. Such a significant instability in the imaging system would produce a noticeable degradation of the non-coronagraphic instrument PSF and, as noted in Sec. 5.2.2, there is also no indication of such displacements to the coronagraphic extinction profile on timescales similar to the 8 s exposure time in raw Ceres data frames: frame-to-frame positional variations of as little as 1/10th of a pixel RMS would already be noticeable by eye. It is therefore highly unlikely that either high-frequency instrumental jitter, or errors in the co-alignment and stacking of the Ceres frames during data reduction, are responsible for the observed discrepancy between the measured and modelled ALC extinction profiles.

5.3.4.4 Discussion and choice of final model

We can conclude from the above discussion that either NCPEs or an out-of-focus ALC mask can individually explain the observed discrepancy with Ceres observations to the precision required for accurate calibration of the coronagraphic system. However, this is only possible if these effects are allowed to significantly exceed their respective SPHERE performance tolerances. This naturally leads to a more likely alternative solution, where a combination of NCPEs and mask defocus can be invoked to accurately reproduce the observed extinction curve whilst remaining consistent with their respective tolerance limits.

However, for the remainder of this investigation we shall adopt the 90 nm RMS (4 mm) ALC focal-plane mask shift error model. Not only does this provide the best fit to on-sky data, but being a single-parameter and time-invariant effect it is also the simplest of the two models, which if correct is applicable to all observations. Since the two effects were found to be more-or-less identical in their effect on the ALC extinction curve, it is not helpful in this instance to construct a complex multi-parameter model which over-fits the available data. Most importantly, with only one on-sky dataset available there is no information on the time-variability of the extinction profile, making it difficult to invoke a NCPE model without first assessing its impact under changing instrument conditions, especially for observations where active NCPE compensation routines are used.

In order to build a more robust model for future calibration efforts, we strongly advise that a measurement of the full coronagraphic extinction profile be made as part of the standard calibration procedure for coronagraphic instruments. Ideally this measurement should be repeated over multiple nights in conjunction with NCPE wavefront sensor measurements, to determine the stability of the extinction profile and the underlying source of any degrading effects.

5.4 Calibrating Protoplanetary Disk Observations

The approach of coronagraphic extinction profile calibration will perform most accurately on polarimetry data, which contain significantly fewer stellar residuals and imaging artefacts than intensity data reduction techniques and is highly favoured for studying protoplanetary disks. When imaging disks using the dual-band polarimetric imaging (DPI) mode of SPHERE-IRDIS, there are however some additional caveats which are not applicable to the Ceres observations discussed up until this point.

5.4.1 Polarimetric image formation and instrumental polarisation

Polarimetric imaging artefacts predominately arise from differences in the image formation process for DPI observations, where the polarised intensity (PI) image is reconstructed based on sets of $Q_\pm$ and $U_\pm$ images, in which a linear polariser is angled at $0/90^\circ$ and $45/135^\circ$ respectively with respect to the image plane. A practical overview of PDI-mode imaging can be found in de Boer et al. (2017), while the core idealised equations for this reconstruction process are summarised below for convenience:

$$I = (Q_+ + Q_- + U_+ + U_-)/2, \tag{5.4}$$

$$Q = Q_+ - Q_-, \tag{5.5}$$

$$U = U_+ - U_-, \text{and} \tag{5.6}$$

$$PI = \sqrt{Q^2 + U^2}. \tag{5.7}$$

The first caveat for DPI observations is that sources of instrumental polarisation, combined with the presence of a bright central source, often create significant spurious signals in PI images, particularly at small angular separations. A common approach to overcome this limitation is to transform this polarisation signal into a polar coordinate system:

$$Q_\phi = Q\cos(2\theta) + U\sin(2\theta) \tag{5.8}$$

$$U_\phi = Q\sin(2\theta) - U\cos(2\theta). \tag{5.9}$$

It is then possible to analyse only the azimuthal polarisation contribution (positive Q_ϕ) since this contains the majority of astrophysical signal, while any residual U_ϕ signal is assumed to be instrumental in nature (Avenhaus et al. 2014b).

A more robust approach to calibrate for this instrumental polarisation is to apply a detailed model of the polarimetric crosstalk properties of the imaging system, in order to directly simulate and remove the residual instrumental polarisation contribution: such a model is provided for VLT-SPHERE by van Holstein (2016) and is applied in Sec. 5.5.

5.4.2 Convolutional depolarisation: an additional consideration for DPI-mode observations

For DPI-mode observations, the coronagraph is also no longer the only source of suppression for linearly polarised disk flux at small inner-working angles. The “butterfly polarisation erasure” effect (Heikamp et al., in prep.) results in an artificially depolarised hole in the center of the image, due to the polarimetric image formation process using a finite PSF.

Following Heikamp et al. (in prep.), the cause and impact of this polarisation erasure effect is illustrated in Fig. 5.8 for a constant-intensity, perfectly azimuthally polarised disk model with no central cavity. Polarimetric differential imaging of such a target in Stokes Q and U produces so-called butterfly patterns (lower panels), which display a characteristic quadrupole signal which tends towards a singularity at the stellar position. When this disk is imaged with a finite instrument PSF, convolutional blurring of the butterfly pattern mixes the positive and negative regions, an effect which is particularly pronounced within a few λ/D of the central singularity. This results in depolarisation and hence loss of disk signal in the total linearly polarised intensity (PI) image, creating the illusion that there is an inner cavity in the disk. The bottom panel of Fig. 5.8 shows that this convolutional polarisation erasure effect (green curve) has a significant impact on the retrieved radial profile out to angular separations as large as 1 arcsec, for the simplest case of a uniform disk of infinite extent. Its contribution to total disk signal loss is also larger than that of the coronagraph alone (blue curve) for separations below 100 mas, with the profile produced by coronagraphic DPI-mode observations (red curve) affected by both sources of extinction. Here the term “un-occulted” is used to refer to observations where the ALC focal-plane mask is not aligned with the target, in the same way as the off-axis Ceres images of Fig. 5.1. It should be noted that the convolutional polarisation erasure effect will be somewhat lower for observing modes which do not use the ALC coronagraphic system, since the Lyot mask and apodiser, which are still included in the green curve of Fig. 5.8, result in a significantly broader PSF and elevated diffraction structure when compared with the PSF of the VLT aperture alone.

The preceding discussion serves to emphasise that attempting to normalise DPI-mode disk observations with a coronagraphic throughput curve derived from intensity

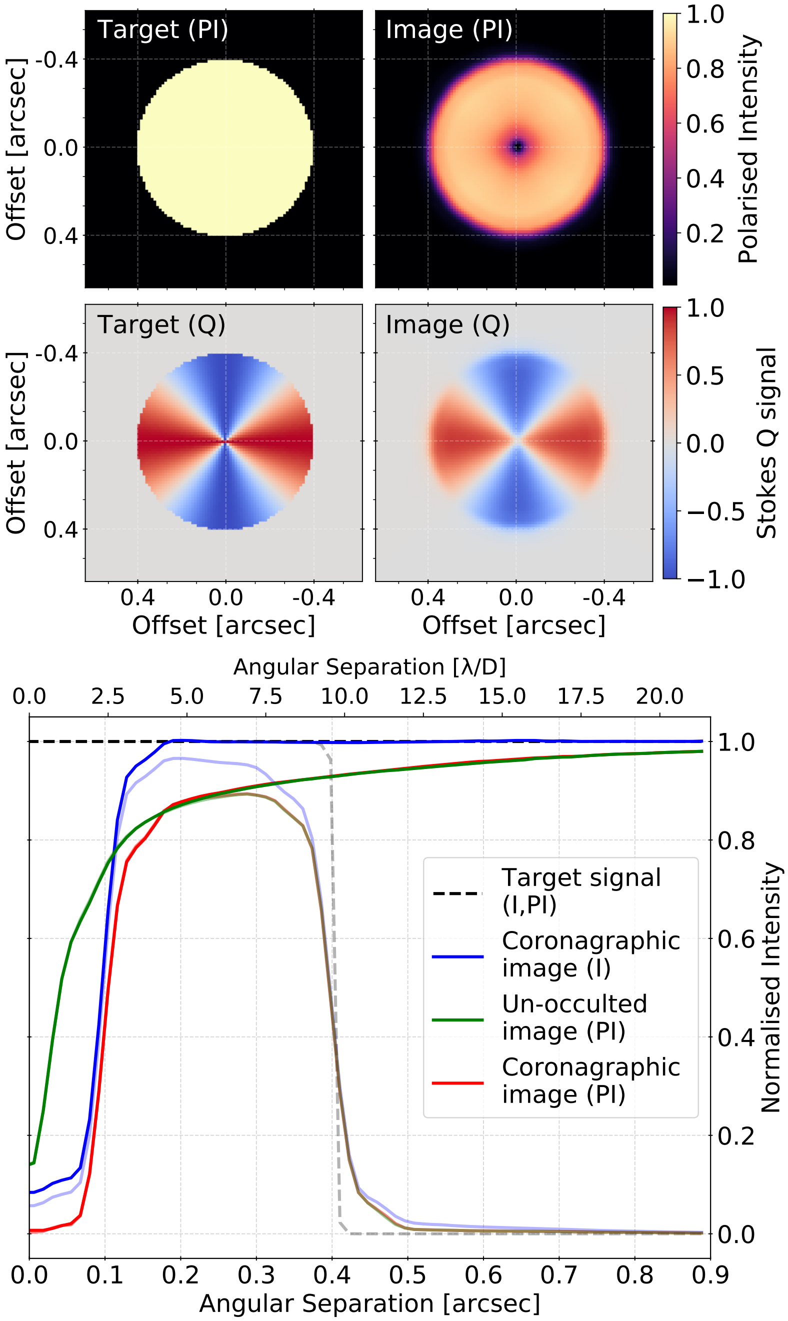

Figure 5.8: Impact of convolutional polarisation erasure on DPI-mode observations, under diffraction-limited VLT observing conditions. *Top panels:* Simulation of the polarimetric image formation process for a finite, uniformly 100% polarised disk target (upper left). The perfect Stokes Q (lower left) and convolved Stokes Q image (lower right) show a quadrupole butterfly pattern, the blurring of which leads to a depolarisation of the innermost regions of the polarised intensity (PI) image (upper right) when combined with the companion Stokes U image. *Bottom panel:* Radial intensity plot of signal extinction contributions for the various imaging modes. Faded lines denote the case of the finite 400 mas radius disk of constant-intensity shown in the upper panels, while opaque lines show the simpler case of a disk of infinite extent. The un-occulted PI image (green) is solely affected by polarisation erasure, while the coronagraphic PI image (red) is affected by both sources of extinction.

imaging alone (such as the red curve of Fig. 5.4) will still result in a systematic underestimate of the true surface brightness of the target. Due to the relative strength of this effect at larger radii, it should be properly accounted for not just when trying to recover signal in the innermost few hundred mas, but whenever there is significant disk signal inside 1 arcsec. Since the effect is inherently source-dependent, the only robust way do this is by forward modelling the disk in question through the raw $Q_\pm$ and $U_\pm$ images in which the effect is produced.

5.4.3 A recipe for disk calibration

The most straightforward approach to calibrating disk observations for coronagraphic extinction is to simply divide out the measured direct throughput of the coronagraphic system as a function of angular separation. Via this method it is in principle possible to correct for the influence of the coronagraph beyond approximately 100 mas to an accuracy of better than 2%, although the location of this radial limit depends somewhat on the flux distribution of the target in question.

As has been shown in Fig. 5.4 and Fig. 5.8 however, this normalisation approach will not perform adequately when the observer is either interested in accurately recovering the relative photometry of angular separations below approximately 100 mas, or is reducing DPI-mode polarimetric observations with significant disk signal interior to approximately 1 arcsecond. These factors, due to diffracted light and convolutional polarisation erasure respectively, are inherently dependent upon the flux distribution of the disk itself. Because of this there is no single profile which is capable of properly calibrating all coronagraphic observations via simple normalisation.

Instead, the most robust approach is to forward-model the disk in question, using the best model of the spatially-variable coronagraphic PSF from Sec. 5.3.4 with the polarimetric image formation procedure of Sec. 5.4.2, and hence properly account for both effects. This approach can be explicitly formulated in terms of the following steps:

1) Identify the centre of the focal-plane mask, which defines the zero-point for coronagraphic extinction calibration. This is best achieved using a flat-field reference image taken with the mask in the beam, but if not available it may be possible to estimate this from the inner regions of the disk image assuming no obvious asymmetric structures are present.

2) Generate coronagraphic instrument PSFs $p_{ij}(x, y)$ via Eq. 5.2 at the appropriate wavelength band centre for each resolution element of interest, with the simulated coronagraphic mask correctly centred with respect to step 1.

3) Identify the stellar position to define the origin for polarisation erasure calibration. This is ideally obtained using a centring frame, where artificial satellite copies of the star are created via the AO system in order to identify the location of the star under the coronagraphic mask.

4) Create an object model, preferably with as few free parameters as possible, which is capable of accurately describing the disk regions of interest. For inclined disks or those possessing complex features a 2-D or full 3-D model is most likely necessary, but for simpler face-on systems a 1-D radial model is sufficient. This is then projected and/or interpolated to create a 2-D model signal map, which constitutes a “perfect” observation of the target.

5) If reducing polarimetry data, create model $Q\pm$ and $U\pm$ frames: this is most simply achieved by assuming perfect azimuthal polarisation of the disk signal with respect to the stellar position from step 3. More realistic polarisation models including factors such as the angular-dependence of polarisation efficiency by scattering (e.g. Pinte et al. 2009) can also be performed in this step, to naturally account for such factors during the reconstruction procedure.

6) Propagate the resulting 2-D model signal map(s) through the imaging system by element-wise convolution with $p_{ij}(x, y)$ according to Eq.s 5.1, and perform Q_ϕ image reconstruction using Eq.s 5.4-5.9 if appropriate.

7) Apply a (preferably global) optimisation routine to derive the most likely disk model for the given observations, passing each iteration through the image formation process of steps 4 - 6.

Deconvolution algorithms (such as the Richardson-Lucy algorithm, Richardson 1972; Lucy 1974) could also in principle be used to invert the problem and directly estimate an object model based on the image data and full instrument model. However, most such algorithms are not capable of operating with a spatially variable PSF and hence coronagraphic observations. They are also prone to generating unphysical artefacts, including “ringing”, due to image noise and small errors in the instrumental PSF model, given the inherent degeneracy of the convolution operation itself (see e.g. Mosleh et al. 2014).

In the following section we re-reduce the DPI-mode observations TW Hydrae presented in van Boekel et al. (2017), in order to illustrate the approach outlined above for calibrating the inner regions of protoplanetary disks.

5.5 TW Hydrae: calibration of a real disk observation

5.5.1 Background and observations

TW Hydrae is the closest T-Tauri star to the sun, located at a distance of $59.5^{+0.96}_{-0.93}$ pc (Gaia Collaboration et al. 2016a). It hosts a nearly face-on (i≈ 7°, Qi et al. 2004) transition disk with a very high degree of radial symmetry, which is particularly gas-rich for its estimated age (3-15 Myr, Sokal et al. 2018). Due to its close proximity to Earth, this object has been the subject of extensive study at infra-red through to sub-mm wavelengths, most notably with SPHERE (e.g. van Boekel et al. 2017, hereafter vB17) and ALMA (e.g. Andrews et al. 2016), the latter at an unprecedented spatial resolution of 1 AU. Scattered light observations reveal three prominent depressions in polarised intensity signal at 1.57” (94 AU), 0.39” (23 AU) and <0.11” (6.6 AU), when scaled with r^2 to correct for the decrease in stellar illumination as a function of radius. These variations have been successfully modelled using self-consistent radiative transfer (RT) code in vB17 as variations in surface density of micron dust grains in the upper layers of the disk. Andrews et al. (2016) reports the detection of depletion features in mm-sized dust grains at similar angular separations, in addition to an inner depletion feature on scales of 1 AU (20 mas), which could be an indication of planet formation: this last feature lies significantly below the 92.5 mas (5.5 AU) mask edge of the SPHERE ALC coronagraph.

In 2015 vB17 observed TW Hydrae in non-coronagraphic DPI mode in the R' and I bands with ZIMPOL, in addition to non-coronagraphic J-band and coronagraphic H-band observations with IRDIS. Here we use the coronagraphic H-band dataset for our analysis, which makes use of the N_ALC_YJH_S coronagraph system addressed in this work. The H-band Q_ϕ image of TW Hydrae is shown in the upper panel of Fig. 5.9, which displays two broad ring features within the central arcsecond and strong coronagraphic extinction in the centre-most region. This image has been corrected for instrumental polarisation following the approach of van Holstein (2016). Of particular interest to this work is the feature termed gap #3, which is located below approximately 110 mas. This is closely coincident with the shoulder of the ALC extinction profile and so provides a useful validation for our calibration approach. vB17 concluded that gap #3 is a real surface density variation: this is supported both by SED studies (Calvet et al. 2002) and their own RT modelling results, which were forward-modelled via $Q_\pm$ and $U_\pm$ to account for the DPI-mode extinction effect of Sec. 5.4.2.

The lower panel of Fig. 5.9 shows the radial intensity curve of this disk (black curve), azimuthally averaged about the disk centre, which is inferred based on the centroid of the ring #2 and ring #3 features outside the ALC region of influence. Since the disk of TW Hydrae is almost directly face-on and displays no significant azimuthal variations, we may readily perform the ALC calibration using a 1-D radial model, without the need for full 3-D modelling of the target.

The following subsections present the results of the two independent calibration approaches outlined in Sec. 5.4.3: both use the final model of the N_ALC_YJH_S coronagraphic imaging system developed in Sec. 5.3, which includes a 4 mm-defocussed focal-plane mask but is otherwise diffraction-limited.

5.5.2 A simple forward model of the TW Hydrae disk

In order to construct an appropriate object model for the disk, it is first important to estimate the likely components of the signal at each angular separation. This is shown by the coloured lines in the lower panel of Fig. 5.9, which simulate the coronagraphic DPI-mode image of a perfectly uniform r^0 intensity profile (similar to that used in Fig. 5.8), in addition to those of $1/r$ and $1/r^2$ intensity profiles. In the idealised case this last profile corresponds to a disk of constant surface density, since stellar illumination of the disk also falls off with $1/r^2$. The solid lines in this plot show the simulated images resulting from coronagraphic DPI-mode observation of each of the three disk profiles, which are scaled to match the most appropriate region of the observed TW Hydrae radial curve. From this it can immediately be concluded that a disk with constant surface density ($1/r^2$ intensity profile, red curve) cannot by itself explain the observed signal in the inner regions, since it does not match well with any part of the bright ring #3 feature interior to 250 mas. A better model for the outer shoulder of ring #3 from 175 to 250 mas is a 1/r radial intensity profile (green curve), however this also over-estimates the signal interior to the 175 mas peak of this feature, and so cannot continue any further in towards the star. A uniform intensity profile (blue curve) can however provide a near-perfect match to the innermost angular separations of the coronagraph, below this 175 mas peak. By contrast, the outer fall-off of ring #3 beyond 250 mas is significantly faster than $1/r^2$, before flattening out at approximately 400 mas.

This first-order analysis of the likely disk structure motivates a multi-component

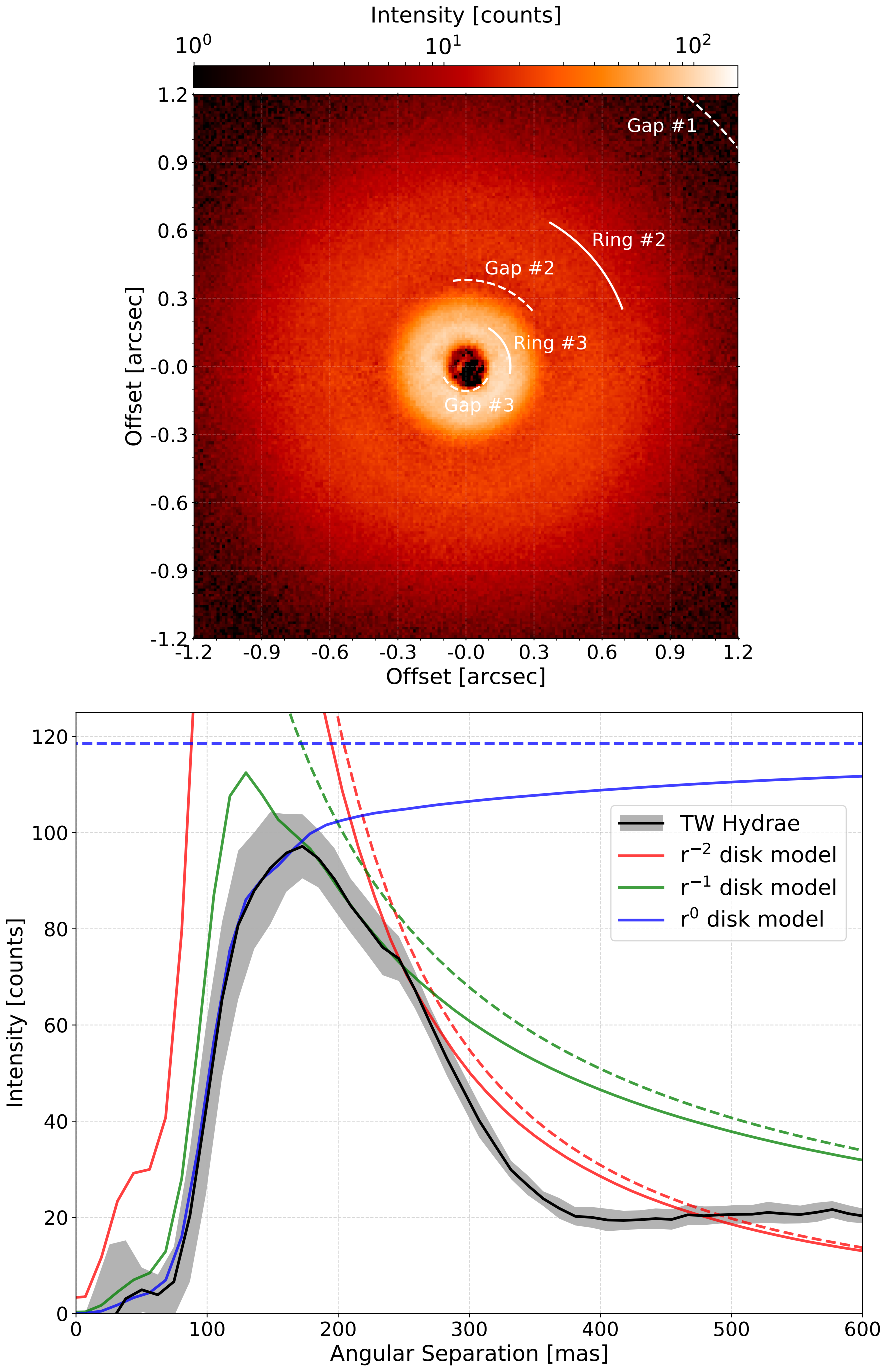

Figure 5.9: H-band coronagraphic Q_ϕ profile of the TW Hydrae disk, after correction for instrumental polarisation via the approach of van Holstein (2016). Intensity is not in this case normalised by r^2, so as to highlight the impact of coronagraphic extinction. ***Top:*** Logarithmically scaled Q_ϕ intensity image, with features labelled according to the convention of van Boekel et al. (2017). ***Bottom:*** Radial intensity plot of the Q_ϕ image (black curve), with the grey shaded region denoting the 1σ bound on azimuthal variability. Three simple disk models are over-plotted, with dashed lines corresponding to the true disk signal and solid lines to the simulated coronagraphic Q_ϕ image for each model, each of which are scaled to match appropriate regions of the TW Hydrae disk.

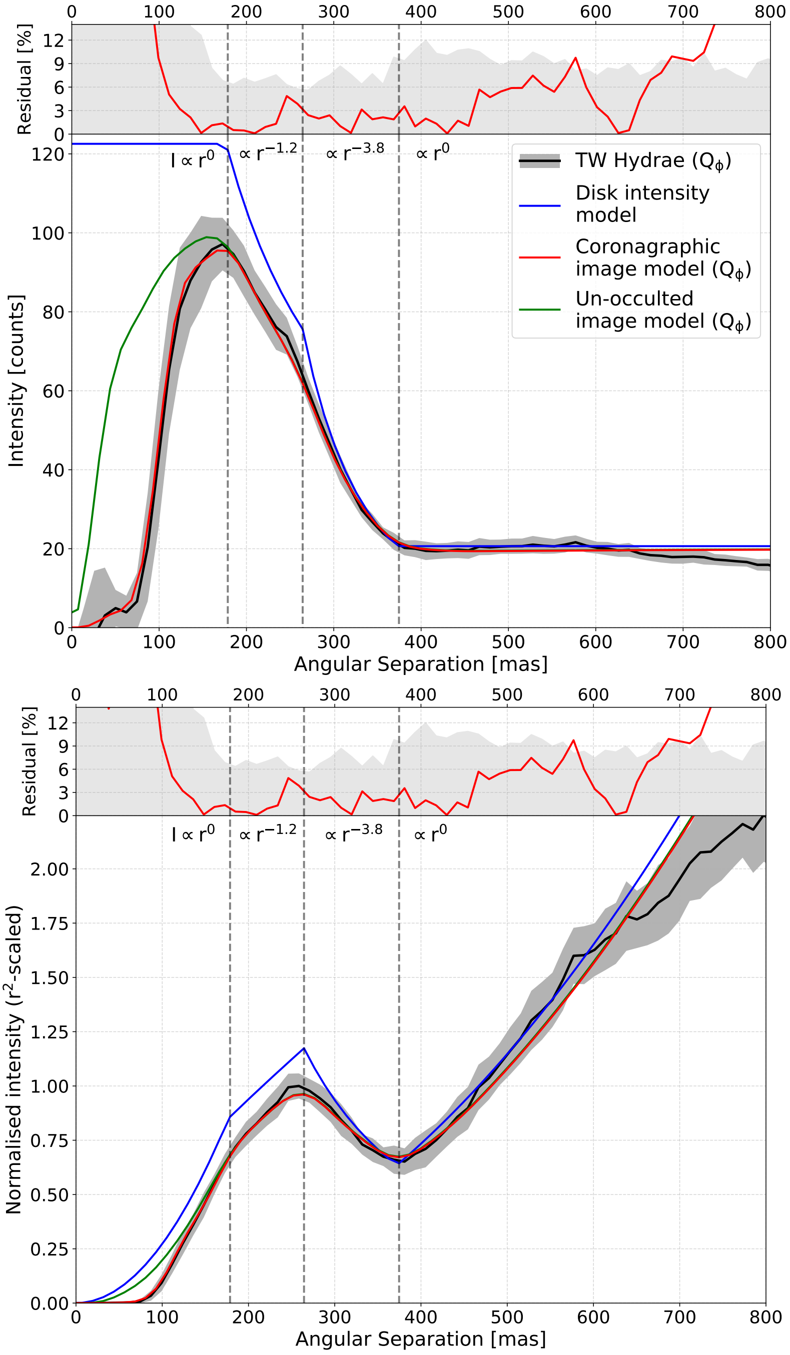

Figure 5.10: Comparison of a simplistic four-power law model of the inner 600 mas of the TW Hydrae disk (blue curve) with vB17 coronagraphic H-band observations (black curve). The upper figure shows this in an un-normalised intensity scaling as in Fig. 5.9, while the lower figure shows the same data scaled with r^2 to correct for the radially decreasing stellar illumination, normalised to unity by the peak of the ring #3 feature. The red and green curves in the main panel show the simulated coronagraphic and un-occulted Q_ϕ images, respectively, which result from the chosen four-power-law model. Percentage residuals between the coronagraphic image model and the observed radial curve are also shown in red in the upper inset panels, where the the grey shaded region denotes the 1σ bound on azimuthal variability of the TW Hydrae disk.

power-law model, with differing spectral index as a function of angular separation. One such model is shown in Fig. 5.10, which is constructed to reproduce all signal interior to 600 mas by using four separate power law components. The intensity profile $I(r)$ of each regime as a function or radius r is described in general by $I(r) = a_1 r^{a_2} + a_3$, where a_1, a_1 and a_3 are free parameters, the final coefficients of which are listed in Table 5.1. This model in fact possesses only three truly free parameters; the two non-zero spectral indices (a_2) and the constant offset (a_3) beyond 375 mas. All other values are constrained either to be zero, or by the requirement of continuity across the three radial boundaries, the locations of which are chosen to correspond with those identified in the previous paragraph based on Fig. 5.9. The resulting model consists of two r^0 regimes for the innermost and outermost regions of the disk, with an approximately $1/r$ profile describing the 175 to 250 mas region and a faster $r^{-3.8}$ describing the subsequent fall-off of the outer edge of the ring feature.

The upper panel of Fig. 5.10 shows that the Q_ϕ coronagraphic simulated image of this simple model (red curve) is capable of providing an excellent fit to the observed disk profile, remaining within the 1σ azimuthal variability limit (grey shaded region) for all separations below 700 mas. Most importantly, this shows that there is no evidence for a decrease in detected disk intensity below angular separations of 180 mas in the Q_ϕ image, once instrumental effects have been corrected for. The equivalent un-occulted Q_ϕ image model (green curve) is identical to the coronagraphic red curve except for the removal of the ALC focal-plane mask, and illustrates how significant the convolutional polarisation effect is for this disk. While the ALC coronagraph begins to cause instrumental signal suppression below 180 mas, noticeable systematic losses due to polarisation erasure already begin to occur at angular separations as large as 300 mas. This outer cutoff is smaller than that quoted in Sec. 5.4.2 since there is relatively little disk signal beyond 300 mas, but it remains arguably a more important effect than coronagraphic extinction for the correct inference of the underlying disk morphology at small angular separations.

The lower panel of Fig. 5.10 plots the same data, but using an r^2 intensity scaling to correct for stellar illumination and hence more fairly present the relative scattering signal from different regions of the disk. In this illumination-corrected scaling, all regions of the model with power-law indices shallower than r^{-2} are seen as radially-increasing, which in turn indicates that for this model there is a relative dearth of scattering signal from the innermost regions of the disk. If this feature may be assumed to originate from variations in the density of micron-sized dust grains in the upper layers of the disk (as opposed to shadowing effects due to surface scale height variations at smaller radii), then this model indicates that there must be a dearth of scattering medium for angular separations smaller than approximately 250 mas: as already shown by Fig. 5.9 it is impossible to construct an accurate model of the observed profile which does not contain such a depression feature from 250 mas down to at least 100 mas. This in turn confirms that the <6.6 AU gap #3 identified by vB17 cannot be purely instrumental in nature.

It should be noted that this forward-modelling process is inherently degenerate at small angular separations due to the high levels of extinction at these locations, and that this proposed model is only one of many possibilities. As may be expected, it was seen that almost any flux distribution underneath the 92.5 mas radius ALC focal-plane mask provides a good match to the observations. For detailed characterisation studies it is therefore recommended that a more thorough forward-modelling approach such

r_0 [mas]	r_1 [mas]	a_1 [counts/mas]	a_2	a_3 [counts]
0	179	0	0	123
179	264	6.10×10^4	-1.2	0
264	375	1.21×10^{11}	-3.8	0
375	∞	0	0	20.6

Table 5.1: Coefficients for the four-power-law model ($I(r) = a_1 r^{a_2} + a_3$) of the inner 600 mas of the TW Hydrae disk. r_0 and r_1 denote the inner and outer boundaries of each regime, respectively.

as a Monte Carlo Markov chain should be used, in order to properly constrain the distribution of possible object models. Such a detailed and computationally expensive analysis is however beyond the scope of this paper.

5.5.3 Assessing the accuracy of ALC extinction calibration by normalisation

The preceding analysis has shown how, by using full model of the spatially variable PSF of the imaging system, it is possible to account for both coronagraphic extinction and convolutional polarisation erasure and determine a simple but accurate model of the underlying disk signal. The final goal of this investigation is to evaluate the relative effectiveness of the simpler, more readily-applicable calibration technique; normalising for the direct throughput of the coronagraphic system. This is significantly faster to implement since it does not require forward modelling as in the previous results, where each step in the computation consists of a computationally expensive element-wise convolution with the spatially-variable instrument PSF. The downsides of this approach are that, as discussed in Sec. 5.4.3, the normalisation calibration does not take polarisation erasure into account and will also over-estimate the innermost regions underneath the ALC mask, where the signal is dominated by non-useful diffracted signal. It is therefore important to determine how far this approach can be trusted for the example of a real protoplanetary disk.

Calibration was achieved by co-aligning a 2-D interpolation of the direct throughput curve model with the ALC mask centre for the TW Hydrae observations, which was itself estimated by minimizing the azimuthal variation of the retrieved radial profile between 80 and 125 mas. The TW Hydrae image was then normalised using this 2-D profile and the radial profile and azimuthal standard deviation computed around the measured disk centre. The normalised image and resulting radial profile are shown in Fig. 5.11: these can be seen to display significant residuals in the inner regions of the disk, which are most likely due to the amplification of low-level instrumental polarisation residuals. Nonetheless, when the median profile of this normalised dataset (solid brown line) is compared with the non-coronagraphic simulated Q_ϕ image from Fig. 5.10 (shown in green in both figures), an extremely close match can be seen down to an angular separation of 75 mas.

Such a close correspondence is in fact expected, since both calibration approaches use the same underlying model of the coronagraphic system and the forward model is tailored to reproduce the coronagraphic TW Hydrae image. It does however clearly illustrate that the simple normalisation approach can be statistically trusted in this case for all angular separations above 75 mas, which in turn implies that there is no

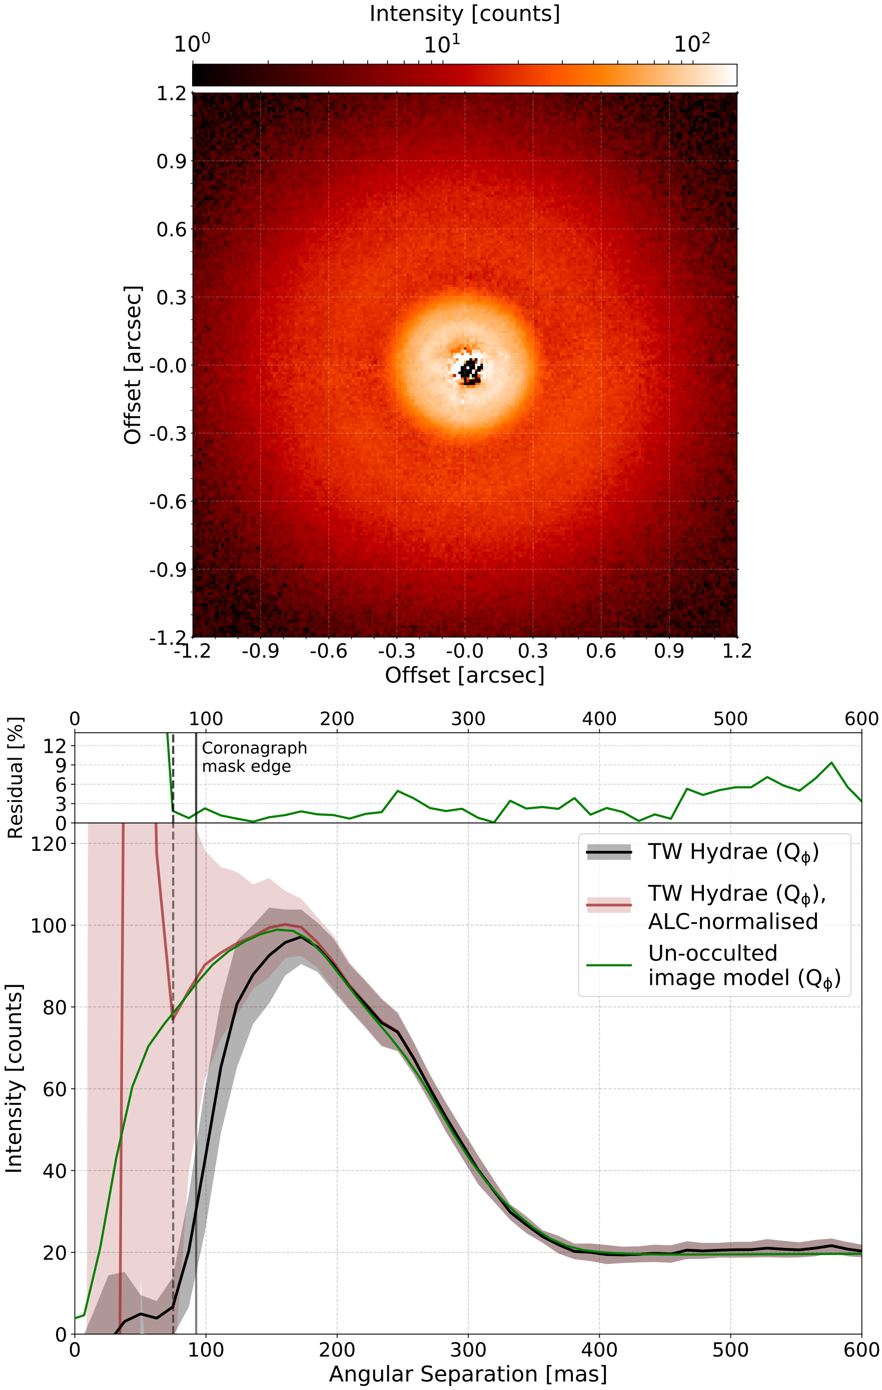

Figure 5.11: Correction of the TW Hydrae Q_ϕ image for coronagraphic extinction, by normalising with the simulated direct-throughput curve of the best-fit ALC system model from Sec. 5.3.4.4. ***Top:*** Q_ϕ image on the same logarithmic scale as Fig. 5.9, normalised with a 2-D interpolation of the ALC extinction curve. ***Bottom:*** Comparison of the radial curve of the normalised TW Hydrae data (brown) with the un-occulted image model simulation of Fig. 5.10 (green). The upper panel shows the percentage residual between these two curves (green), with the left-most dashed line denoting the point of divergence at 75 mas. Brown shading denotes the 1σ bound on azimuthal variability of the ALC-corrected TW Hydrae disk.

significant contribution of non-useful diffracted signal at these separations. Such a small inner limit on the calibration accuracy, which lies below the ALC mask edge, is notably lower than the general 100 mas limit derived in Sec. 5.3 based on Fig. 5.4. This most likely stems from the fact that the TW Hydrae disk displays significantly less extended signal at wider separations than the 300 mas-radius, uniform-intensity Ceres disk, resulting in a correspondingly lower diffracted signal contribution from these resolution elements. This is true of most protoplanetary disk observations, and so one may reasonably expect the normalisation approach to be capable of accurately calibrating for coronagraphic extinction down to the 92.5 mas edge of the ALC mask, unless large regions of high surface brightness are observed within the central few hundred milli-arcseconds of the disk.

Despite the fact that the median profile is seen to be accurately corrected down to this 75 mas limit, such small angular separations are however of limited value in the case of TW Hydrae: it is difficult to draw strong conclusions about the radial disk structure below approximately 175 mas based on the brown curve alone, due to the explosion in relative uncertainty that comes with rapidly decreasing ALC throughput. This is of course notwithstanding the systematic modification made to the apparent disk signal by polarisation erasure: in this case one could erroneously claim that, rather than a uniform-intensity (r^0) signal interior to 180 mas, there is in fact evidence for a further central depression in the un-normalised Q_ϕ image.

It is also important to note that these two ALC-corrected model profiles are still not directly comparable to the flux distribution which would be observed via equivalent non-coronagraphic DPI-mode observations. The subtle difference between the two is that non-coronagraphic observation modes do not include the ALC apodiser and Lyot stop, whereas un-occulted (or ALC extinction-corrected) simulations still include the impact of these pupil-plane optics on the instrument PSF. This is crucial when considering the impact of convolutional polarisation erasure, since the effect is significantly worse in the latter case due to the broader instrument PSF. The simple solution for this when performing forward modelling of the target is to use an equivalent optical model in which these two pupil-plane optics have been removed; in this case leaving only the VLT aperture mask. No such solution is however available for the normalisation calibration, which makes it challenging to directly compare datasets taken in different observing modes even after calibrating for coronagraphic extinction.

The authors therefore strongly recommend that all studies using DPI-mode observations should perform a full analysis of the observed disk signal by forward-modelling through $Q_\pm$ and $U_\pm$ frames, in order to take into account the highly non-linear effects of convolutional polarisation erasure. For intensity-mode observations it should in most cases be sufficient to simply normalise for coronagraphic extinction, provided that the contribution of non-useful diffracted signal is shown to be negligible compared to the direct throughput of the coronagraphic system at the angular separations of interest.

5.6 Conclusions

In this work we have demonstrated that observing solar system targets such as Ceres, which are small enough on which to lock the adaptive optics system of SPHERE but extended enough to cover the full radial extent of the coronagraphic extinction profile, provide a useful method for calibrating the coronagraphic imaging system. This approach provides the full two-dimensional extinction profile of the coronagraph and is shown to be insensitive to changes in seeing conditions and upstream wavefront aberrations, but with the caveat that it inherently combines the direct (useful) throughput of the coronagraph with the (non-useful) diffracted signal contribution of the extended target. This makes it highly complementary to the more common point-source calibration approach, which measures the one-dimensional direct throughput profile of the system and is more sensitive to temporal instrumental or atmospheric variations.

When such calibration measurements are used to inform instrument simulation efforts, the resulting model of the spatially-variable point-spread function (PSF) of the instrument provides a powerful tool for properly calibrating observations of circumstellar material. We conclude that in the case of the SPHERE-IRDIS apodised Lyot coronagraph (ALC), the observed extinction profile is best reproduced by invoking a 4 mm longitudinal shift of the ALC mask from the coronagraphic focal plane, in an otherwise diffraction-limited system. It is important to determine whether this is a feasible scenario, or whether other factors such as non-common path error (NCPE) are in fact responsible for the minor observed degradation in coronagraphic efficiency. Most importantly, if a time-variable effect such as NCPE is found to be responsible, a more detailed instrument model would need to be developed to accurately calibrate future observations.

For dual-band polarimetric imaging (DPI) observations, the loss of polarised intensity signal due to convolutional polarisation erasure is seen to be of arguably greater importance than coronagraphic extinction, especially at angular separations larger than 100 mas. The effect is however less severe in non-coronagraphic imaging modes, since it is the broadened instrument PSF produced by the use of apodising pupil-plane optics during ALC observations that significantly worsens the convolutional properties of the imaging system. Due to the highly source-dependent and nonlinear nature of this polarisation erasure effect, it can only be effectively accounted for via appropriate forward modelling of the surface brightness distribution of the target through the raw $Q_{\pm}$ and $U_{\pm}$ frames in which the effect is produced.

The application of this forward-modelling calibration approach to coronagraphic DPI-mode observations of TW Hydrae has shown that it is possible to adequately explain all observed signal from the inner 600 mas of the disk using a four-component one-dimensional power-law model. The innermost signal depression feature (gap #3) observed by vB17 is in this way confirmed to be a real property of the disk, and not due to artefacts of the coronagraphic or polarimetric imaging process. The simpler calibration approach of direct-throughput normalisation is also capable of self-consistently correcting the median radial profile of the disk for coronagraphic extinction down to an angular separation of 75 mas, significantly below the 92.5 mas edge of the ALC focal-plane mask. While this makes the approach suitable for the calibrating the majority of intensity-mode disk observations, its major limitation for DPI mode is that it cannot be used to correct for polarisation erasure effects, which also prevents a direct comparison with non-coronagraphic observations. Finally, it was seen for both calibration approaches that the choice of disk model remained largely unconstrained below the 100 mas inner-working-angle of the coronagraph, due to the rapid increase in fractional uncertainty with decreasing system throughput at smaller angular separations.

Such a calibration approach is nonetheless indispensable when attempting to correctly interpret the relative surface brightness profiles of the innermost regions of protoplanetary disks, especially between angular separations of approximately 100 to 300 mas. The approach presented in this paper is applicable not only to the SPHERE-IRDIS ALC system, but equally to any other current or future high-contrast imaging instruments which make use of attenuating coronagraphic optics. This is particularly true for JWST-NIRCam, which will incorporate asymmetric wedge coronagraphs (Krist et al. 2009) which cannot be calibrated using conventional 1-D radial extinction profile approaches. Such a detailed calibration effort will ultimately allow these instruments to maximise their science yield at the smallest possible angular separations, where the discovery space is as-yet largely unexplored.