



HAL
open science

Contrast-based tolerancing of space telescopes for exoEarth imaging

Iva Laginja

► **To cite this version:**

Iva Laginja. Contrast-based tolerancing of space telescopes for exoEarth imaging. Astrophysics [astro-ph]. Observatoire de Paris, Université Paris Sciences et Lettres (PSL); Office National d'Études et de Recherches Aérospatiales (ONERA); Space Telescope Science Institute (STScI); Laboratoire d'Astrophysique de Marseille (LAM), 2021. English. NNT : . tel-03546315v2

HAL Id: tel-03546315

<https://theses.hal.science/tel-03546315v2>

Submitted on 9 Mar 2022 (v2), last revised 14 Nov 2022 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



THÈSE DE DOCTORAT
DE L'UNIVERSITÉ PSL

Préparée au LAM/ONERA/STScI

Financée par ONERA/STScI

**Contrast-based tolerancing of space telescopes
for exoEarth imaging**

Soutenue par

Iva LAGINJA

Le 15 décembre 2021

École doctorale n°127

**Astronomie et Astrophysique
d'Île-de-France**

Spécialité

Astronomie et Astrophysique

Composition du jury :

Anthony Boccaletti Directeur de Recherche, LESIA	<i>Président</i>
Philip Hinz LAO Director, UC Santa Cruz	<i>Rapporteur</i>
Christophe Vérinaud Adaptive Optics Physicist, ESO	<i>Rapporteur</i>
Rebecca Jensen-Clem Assistant Professor, UC Santa Cruz	<i>Examineur</i>
Raphaël Galicher Maître de Conférence, LESIA	<i>Examineur</i>
Vanessa Bailey Instrument Technologist, NASA JPL	<i>Examineur</i>
Rémi Soummer Associate Astronomer, STScI	<i>Directeur de thèse</i>
Laurent Mugnier Directeur de Recherche, ONERA	<i>Directeur de thèse</i>
Jean-François Sauvage Maître de Recherche, ONERA/LAM	<i>Directeur de thèse (invité)</i>

Contrast-based tolerancing of space telescopes for exoEarth imaging

Dissertation

submitted in partial satisfaction of the requirements for the degree of

Doctor of Philosophy
in
Astronomy and Astrophysics

defended on Wednesday 15 December 2021
at 16:00 hours in Paris, France

by

Iva Laginja

born in Zagreb, Croatia
in 1992

Tolérancement des télescopes spatiaux fondé
sur le contraste pour l'imagerie d'exoTerres

“Gledamo zvezdice”

CONTENTS

1	Introduction	1
1.1	The search for new worlds	1
1.1.1	History and methods of exoplanet detection	1
1.1.2	The challenges of direct imaging	2
1.1.3	State of the field of high-contrast imaging science	4
1.2	Current and future landscape of high-contrast imaging facilities	10
1.2.1	Ground-based observatories	11
1.2.2	Space-based observatories	13
1.2.3	The role of high-contrast imaging testbeds	16
1.3	High-contrast imaging in theory and implementation	19
1.3.1	Image formation and coronagraphy	20
1.3.2	Wavefront sensing and control	26
1.3.3	Post-processing	30
1.4	Application of HCI to the case of large segmented telescopes	31
1.4.1	Wavefront errors and cophasing of segments	32
1.4.2	Ultra-stable telescopes	33
1.5	This thesis	34
2	A method for analytical tolerancing of segmented telescopes	37
2.1	Introduction	40
2.2	PASTIS model of telescope segment-level aberrations in high-contrast coronagraphy	42
2.2.1	Matrix formalism to calculate the average dark hole contrast	42
2.2.2	Semi-analytical calculation of the PASTIS matrix	45
2.2.3	Validating the semi-analytical PASTIS matrix	46
2.3	Model inversion and statistical mean contrast derivation	47
2.3.1	Eigendecomposition of the PASTIS model and mode-segment relationship	47
2.3.2	Contrast as a function of the eigenmodes	48
2.3.3	Statistical mean of the contrast from mode amplitudes	53
2.4	Segment-level tolerance statistics	54
2.4.1	Statistical mean contrast and its variance in segmented coronagraphy	54
2.4.2	Uncorrelated segment-level requirements	55
2.4.3	Case of correlated segments	56
2.5	Application to LUVOIR WFE tolerancing	57
2.5.1	Segment requirements and Monte-Carlo simulations for three APLC designs	57
2.5.2	Modal analysis of the segment-based requirements	59
2.6	Discussion	60
2.7	Conclusions	63
	Appendices	64
2.A	Modal sensitivity analysis on continuous DMs	64
2.A.1	Formalism for continuous DMs	65
2.A.2	Simulated results on HiCAT and RST	66

3	Simulating tolerancing analyses of segmented aperture coronagraphs	71
3.1	A numerical tool for statistical WFE tolerancing	73
3.1.1	Implemented telescopes, coronagraphs and aberration modes	74
3.1.2	Statistical sensitivity analysis	76
3.2	Contrast sensitivities to segment aberrations on a JWST coronagraph	76
3.2.1	Coronagraphy with NIRCcam	78
3.2.2	Sensitivity analysis on NIRCcam with a CLC	78
3.3	Contrast sensitivities to segment aberrations on the HiCAT testbed	83
3.3.1	The HiCAT project and experiment emulation setup	84
3.3.2	Results on the HiCAT emulator	86
3.4	Conclusions	92
4	Wavefront stability tolerancing in experimental validations	95
4.1	Introduction	98
4.2	PASTIS tolerancing model and extension to a drifting coronagraph floor	99
4.3	The HiCAT project and experimental setup	101
4.3.1	Classical Lyot coronagraph as static setup	102
4.3.2	Active wavefront sensing and control for an improved DH	103
4.4	Experimental validation of segmented tolerancing on the HiCAT testbed	105
4.4.1	PASTIS matrix measurement and deterministic forward model validation	105
4.4.2	Validation of mode contrast allocation	108
4.4.3	Statistical validation of independent segment tolerances	111
4.5	Discussion	116
4.6	Conclusions	118
5	Applications of analytical tolerancing in studies for ultra-stable telescopes	121
5.1	Tolerancing approaches for segmented telescopes	123
5.2	Extended spatial tolerancing bases	124
5.2.1	Global modes for low and high spatial frequencies	126
5.2.2	Local Zernike modes	127
5.2.3	Custom local modes from thermo-mechanical modeling	127
5.3	Temporal domain analysis	130
5.3.1	Dynamic wavefront errors and relevant timescales	130
5.3.2	Drift requirements in a closed-loop WFS&C system	136
5.3.3	Perspectives for dynamic WFE tolerancing	137
5.4	Conclusions	139
6	Conclusions and outlook	143
A	Cophasing of a segmented telescope testbed: JOST	147
A.1	Introduction	150
A.2	Testbed description	151
A.2.1	Key hardware components	151
A.2.2	Software upgrades	153
A.2.3	Previous wavefront sensing and control activities	154
A.3	Wide-field wavefront sensing with a segmented deformable mirror	154
A.3.1	Goals of wide-field WFS demonstration	154
A.3.2	Data acquisition and reduction	155
A.3.3	Wavefront sensing results	155
A.3.4	Comparison to previous results without the segmented deformable mirror	156
A.4	Summary and Conclusions	157

B System-level demonstrations of high contrast on a segmented aperture testbed: HiCAT	161
B.1 History and purpose	163
B.2 Testbed optical layout and hardware	163
B.3 Operations and testbed modes	164
Bibliography	169
Acronyms	189
English summary	195
Résumé en français	199
Deutsche Zusammenfassung	205
Hrvatski sažetak	211
Publications and activities beyond research	215
Curriculum vitae	221
Acknowledgments	225



INTRODUCTION

1.1 The search for new worlds

1.1.1 History and methods of exoplanet detection

Among the many scientific problems humanity is exploring today, the question as to whether there is life elsewhere in the universe has been on our minds for centuries. Initially treated mostly under theological and philosophical aspects, the existence of planets around stars other than our own Sun was contemplated by ancient thinkers like Plato and Epicurus (2nd - 5th century BC). It was not until the late 19th and early 20th century though that the first scientifically founded theories about these so-called “exoplanets” were formed and published in peer-reviewed literature. While the first exoplanet detections were falsely claimed in the early 1940s, the research field of observational exoplanetology found its beginnings with the initial discovery of an exoplanet in orbit around a solar-type star in 1995 (Mayor & Queloz 1995), the planet *51 Pegasi b*. The significance of this discovery was underlined in 2019, when Michel Mayor and Didier Queloz won the Nobel prize in physics, implicitly showing recognition to the entire field of exoplanets, having clearly grown out of its infant days of a niche subject in the past three decades.

The deep scientific interest in exoworlds, and the question about the existence of other life forms than our own is confronted with a high complexity in the instrumental methods that let us observe distant bodies. Astronomical instrumentation for exoplanet detection is a multi-faceted field covering an long list of optical systems and techniques, all aiming to capture signs of a potentially habitable planet. Among them, a distinction can be made between *indirect* and *direct* exoplanet detection methods, depending on whether we observe a signal that lets us only infer the existence of a planet, or whether we can capture and analyze the photons from a planet directly. With the ultimate goal to find exoplanets that can actually harbor life, developing these methods to the point where they are sensitive enough to signals from planets that resemble our own Earth, so-called “exoEarths”, is the broad scientific context in which this thesis is placed. After briefly highlighting the most important detection methods used in the exoplanet community, I will take a deeper dive into the only direct detection technique, around which the core of this thesis revolves: direct imaging.

The *radial velocity* (RV) observation technique exploits the fact that a planet in orbit around a star forces both bodies into a Keplerian motion around the system’s barycenter. This means that the star itself will exert relative motions that are indicative of a secondary body around it. The line-of-sight motion of the star for example, its radial velocity, can be observed with spectroscopic measurements that detect the Doppler-shifted absorption lines in the stellar spectrum. This observable can hence be exploited to determine the presence of an orbiting companion. The amplitude of the radial velocity curve is a function of planet mass and orbital inclination, commonly expressed as $M_p \sin(i)$, with M_p the planet mass and i the inclination. An important characteristic of RV is that it is biased toward planets in edge-on orbits because they maximize the observed amplitudes. Additionally, this method is degenerate with respect to mass and inclination of a system and can therefore provide only a lower bound to the exoplanet mass. Considering these two points, we need to use other techniques to better constrain the planet properties.

While the line-of-sight motion of a star can be detected by the change in wavelength in the stellar spectrum, the star’s motion in the plane of the sky can be determined through *astrometry*. This technique

measures the star’s exact position in the sky with respect to a spatial reference frame. Since it visually follows the location of the star, we can measure the system’s orbital inclination, which eliminates the unknown we encounter with the RV method. The major constraint of both of these methods is that they yield a planet mass, but not a radius, which makes them insufficient for any conclusions about the planet’s density, from which we could determine its rough composition.

A method that relies on a different physical phenomenon for exoplanet detection and characterization is planetary *transit photometry*. During such an event, the planet passes in front of its host star, and the disk of the planet covers a part of the stellar surface, dimming the light we observe with photometric measurements. The result is a distinct “dip” in the observed light curve that allows for the determination of the planet’s radius. The first confirmed transiting planet was observed in 1999 (Henry et al. 1999, 2000; Charbonneau et al. 2000), and a plethora of them followed with dedicated transit surveys like conducted with the Kepler Space Telescope. One limitation of this method is that the alignment of the exoplanetary orbit needs to have an inclination such that we actually see the planet travel across the stellar surface, which requires almost perfect edge-on orientations as seen from the Earth. Since a part of the stellar light will be filtered by the atmospheric layers of the planet while it is in transit, we can make spectroscopic measurements of its absorption lines for the purpose of atmospheric characterizations. Similar data can be acquired during a secondary eclipse when the planet is passing behind the star.

Of the ~4800 confirmed exoplanets today¹, roughly 1000 have been detected by radial velocity measurements, and ~3400 by transits. A much smaller subset of detections have been made with other methods, for example *gravitational microlensing*, where we can infer the presence of a planet by the deflection it inflicts on the light of a background source. Noticeably, each observation method can fill a different part of the parameter space for exoplanet characterizations, which is why follow-ups with a method different than the one that was used for detection are needed in order to get a fuller picture of an observed body.

One thing that all the aforementioned detection methods have in common is that they all rely on measurements of the light from the host star, or a different background source, rather than capturing photons emitted from the planet alone. The capability of imaging an exoplanet as an unresolved point source is commonly referred to as *direct imaging*. The planet light captured in this way can either come from its own thermal emission, or as reflected light from its host star. Accessing the light from these extra-solar systems directly enables a generalized spectral study of the planetary atmosphere’s chemical composition and search for the presence of *biomarkers*, our current best guess for indicators of life.

1.1.2 The challenges of direct imaging

To achieve the ambitious goal of direct imaging of Earth-like exoplanets, two fundamental obstacles need to be overcome, the first being the need to resolve the planet as a point source at very small angular separations from its host star. The fundamental *diffraction limit* sets the smallest attainable angular resolution element θ in an aberration-free system by:

$$\theta = \frac{\lambda}{D}, \quad (1.1)$$

where θ is the angular resolution in radians, λ the observing wavelength and D the diameter of the telescope. If we assume an exoplanet at 1 au from its host star at a distance of 10 pc from Earth, the required resolution of 0.1 arcsec lies within the capabilities of existing telescopes, for example the 8 m primary mirrors of the Very Large Telescope (VLT) on the ground or the 2.5 m primary of the Hubble Space Telescope (HST) in space when observed at a wavelength of 700 nm. However, the angular resolution requirement alone, as it turns out, is not the main issue. The second obstacle in direct imaging is the large *planet-to-star flux ratio*, which is more challenging to overcome the closer the planet is to the star, as projected in the sky. Figure 1.1 compares the flux ratios of various planets at different wavelengths in the optical and infrared (IR) range of a solar system analogue at a distance of 10 pc, in units of contrast. Not only do the planets become more luminous at longer wavelengths where they are self-luminous, but the sun itself has less flux in that part of the spectrum, which alleviates the contrast requirement for direct imaging. In the optical range around 700 nm, the planet flux consists mostly of reflected starlight, and the brightness contrast

¹on 20 September 2021, as published by exoplanet.eu




Figure 1.1: Flux ratios for various planetary bodies compared to our Sun, a G-type main sequence star. Displayed are the rocky planets Mars and Earth, the gas-giant Jupiter, and the hot Jupiter exoplanet β Pictoris b . The dashed vertical line at $2 \mu\text{m}$ indicates the observation wavelength for current instruments, for example GPI, where the flux ratio for the hot Jupiter is about four decades. The vertical at 700 nm indicates HCI observations of future space instruments (RST and LUVOIR), where the flux ratio is significantly larger, on the order of 10^{-10} for Earth-like planets. Adapted from Mazoyer (2014).

between an Earth-like planet and a solar-type star is on the order of 10^{-10} . Direct imaging techniques that aim to overcome these brightness ratios fall under the term called *high-contrast imaging (HCI)*.

The need to overcome a large flux ratio is not unique to the detection of exoplanets and had been recognized as one of the main obstacles in observing other faint astronomical sources, for example the solar corona, which is a million times fainter than the surface of the Sun. Until the early 1930s, observing this faint outer component of the Sun has only been possible during solar eclipses when the disk of the moon would move in front of the Sun, suppressing the bright light radiating from its surface, which revealed the fine coronal structures. There were several attempts to image the solar corona outside of solar eclipses in the late 18th and early 19th century, but the first successful observation was done by the French astronomer Bernard Lyot from the Observatory of the Pic du Midi in the southwest of France in 1930. He mounted an occulting disk over the image of the star in the focal plane of the telescope and identified the need for an undersized diaphragm in the reimaged entrance pupil to cut off the light diffracted by the focal-plane mask (Lyot 1939). Thus, the first *coronagraph* (fr.: “coronographe”, Lyot (1932)) was conceived, the Lyot coronagraph, which builds the basis for a whole family of coronagraphs today.

This concept was used in the following decades for observations of the solar corona, to image the outer solar system planets and their rings and satellites, and eventually it was adapted to search circumstellar environments for protoplanetary disks (Vilas & Smith 1987). Using a Lyot coronagraph fabricated by the planetary imaging group at University of Arizona, the first directly imaged circumstellar disk was observed by Smith & Terrile in 1984 (Smith & Terrile 1984) around the 4th magnitude star β Pictoris, visible from the southern hemisphere. It was not until the mid-2000s that coronagraphy was able to reveal the first directly imaged exoplanets (Chauvin et al. 2004), almost a decade after the very first exoplanet discovery in 1995.

High-contrast observations at small angular separations remain a challenging task to this day. Coronagraph designs are advancing to provide a better performance in terms of on-axis light attenuation, while trying to limit the degradation of the planet throughput, all the while moving the lower limit on observable star-to-planet distance closer and closer to the star. As the science goals become more demanding, aiming

to capture the light of very faint, close-in exoplanets, the modern-day *high-contrast imaging instrument* has evolved into more than just a set of static optical masks for on-sky starlight attenuation. Today, a HCI instrument materializes as an integrated system that combines three main components: static starlight attenuation with a coronagraph, active correction of the incoming light with a wavefront sensing and control system (WFS&C), and the application of advanced post-processing algorithms. While we have a whole collection of coronagraph architectures at our disposal, they are all united by their high sensitivity to wavefront errors (WFEs) due to influences external and internal to the telescope. A WFS&C system has the ability to sense the aberrations of the electric field and apply corrections with active components like deformable mirrors (DMs) in order to eliminate aberrations, and to improve the contrast in the final images. In addition to these static and dynamic systems during data collection, the use of advanced post-processing algorithms completes a high-contrast imaging (HCI) instrument with which exoplanet hunters set out to find faraway worlds. A look on the past and current science and HCI capabilities from the ground and from space will describe the technology gap we need to fill in the future if we want to reach the ultimate goal of getting a direct image of an Earth-like exoplanet.

1.1.3 State of the field of high-contrast imaging science

High-contrast imaging aims to reveal the faintest objects as visible both from the ground as well as from space. This usually encompasses faint point sources like exoplanets and brown dwarfs, as well as extended objects with low surface brightness, like protoplanetary or debris disks in circumstellar environments. Ground-based and space-based telescopes complement each other in these observations: ground-based telescopes have a rich history in the development of HCI instruments, with many different coronagraphs being used routinely at facilities around the world. However, atmospheric absorption limits our access to the sky to only certain wavelength windows, which is mitigated by space-based observatories that avoid the atmosphere altogether. Space telescopes have enabled a plethora of new exoplanet discoveries through transit detections, many of which were followed up successfully with through RV from the ground, but only a few direct imaging detections of exoplanets have been made from space. The only satellite currently carrying instruments with sufficient HCI capability is the Hubble Space Telescope, and even there, pertinent observations were possible only with the continuous optimization of the data reduction and post-processing techniques used.

Results from direct imaging observations and surveys

A very fruitful research field for space-based HCI are circumstellar disks (Schneider et al. 2014): deep visible light images like those attainable with satellites above the atmosphere are able to reveal sub-micron dust grains of a stellar system halo. This is much less accessible with ground-based instruments that are more optimized for the near-infrared wavelength regime, which limits observations to micron-size grains in the main planetesimal belt (Schneider et al. 2018). Circumstellar disks thus pose a prime example of complementary research being performed by combining observations from space and from the ground - some examples are displayed in Fig. 1.2. A high degree of attention has of course been given to observations of individual directly imaged exoplanets. Among the most prominent examples, Fig. 1.3 shows images of the stars HR 8799, β Pictoris and 51 Eridani and their imaged exoplanet companions.

For more than a decade after its discovery by Marois et al. (2008), HR 8799 was the home to the only directly imaged multiple-planet system, until the discovery of two giant planets on wide orbits in 2020 (Bohn et al. 2020). The initial discovery of the planets *b*, *c* and *d* was followed by the finding of the closer-in planet *e* (Marois et al. 2010b). The four companions travel on nearly face-on orbits at distances of 10-70 au, and are embedded between a warm inner, and a cold outer debris disk (Matthews et al. 2013; Su et al. 2009). This resembles an upscaled model of our own solar system with the Asteroid belt on the inside of the giant gas planets, and the Kupier belt beyond. Due to its uniqueness, the system has been extensively studied over the past decade, with a high interest in the study of the planets' orbital motions. Improving data reduction and post-processing techniques proved crucial for this goal: Lafrenière et al. (2009) went back to HST data from 1998 and were able to recover the planet HR 8799 *b* by deploying the locally optimized combination of images (LOCI) algorithm (Lafrenière et al. 2007) to construct optimized




Figure 1.2: HCI observations of circumstellar environments reveal structures that let us study the properties of stellar systems. *Top:* Edge-on debris disk around AU Microscopii imaged with the IRDIS subsystem on VLT/SPHERE (Boccaletti et al. 2015b). *Bottom left:* Debris disk around Fomalhaut, imaged with the STIS instrument on HST (Kalas et al. 2013). The circle marks the position of the star and the diamond marks the geometrical center of the dust belt. *Bottom right:* The young star AB Aurigae and the spirals of its protoplanetary disk, imaged in polarized light with VLT/SPHERE (Boccaletti et al. 2020b).




Figure 1.3: Directly imaged exoplanets in their stellar systems. *Left:* HR 8799 with four co-orbiting planets *b*, *c*, *d* and *e*, imaged with Keck (Marois et al. 2010b). *Middle:* β Pictoris with the hot Jupiter β Pic *b*, observed with the VLT (Lagrange et al. 2010). *Right:* 51 Eridani with a young, Jupiter-like companion, detected with GPI (Macintosh et al. 2015).




Figure 1.4: Orbital fitting of the four planets in the HR 8799 system with different dynamical constraints. The black dots show measured astrometry of the exoplanets circling their host star, indicated with the black star in the middle of each plot. 200 orbits are plotted for each of the three cases: unconstrained fitting of Keplerian orbits (*left*), 1:2:4:8 near-coplanar orbits with low eccentricity (*middle*), and exclusively dynamically stable coplanar orbits (*right*). Source: Wang et al. (2018).

reference images. Soummer et al. (2011) optimized the process with LOCI even further by using improved calibration files and recovered not only planet *b*, but also *c* and *d* in the HR 8799 system from archival data. Being able to retrieve orbital measurements from epochs more than a decade apart is essential to perform dynamical studies of these very long-period planets (50-500 yr), and to constrain their masses (Wang et al. 2018). Figure 1.4 displays several results of orbital fitting on astrometry with milliarcsecond precision of the HR 8799 system from data taken by the ground-based Gemini Planet Imager (GPI), showing how different dynamical constraints in orbit fits influence the mass estimates of the planets. In this case, more constrained orbits favor lower planet masses, $\sim 8 M_{Jup}$ for planets *c*, *d* and *e*, and $\sim 6 M_{Jup}$ for planet *b*.

β Pictoris is a young (~ 20 Myr, Mamajek & Bell (2014); Y Navascués et al. (1999)), 4th magnitude star in the southern hemisphere. It started gathering attention after it had been identified as the first star hosting an edge-on circumstellar disk (Smith & Terrile 1984). Closer inspection of the disk showed a deformation in it (Lagage & Pantin 1994) and among the possible explanations was the hypothesis this warp might be caused by the gravitational influence of an exoplanet (Mouillet et al. 1997; Roques et al. 1994). In 2009, Lagrange et al. (2009) reported a point-like signal in observations with the NaCo instrument on the VLT dating back to 2003, which was later confirmed to be the planet β Pic *b* (Lagrange et al. 2010). After it was seen that β Pic *b* has an edge-on orbit, one of the main points of attraction toward this $11 M_{Jup}$ planet was the possibility of seeing the first directly imaged planet in transit. This hypothesis was ruled out for the planet itself (Wang et al. 2016), as it was predicted to miss the star in the end, but its Hill sphere would still travel across the star, providing the opportunity to probe circumplanetary material during a Hill sphere transit (Kenworthy et al. 2021).

The star 51 Eridani has very similar properties to β Pictoris as they are both members of the same moving group. The 2–10 M_{Jup} planet 51 Eri *b* is thus a similarly young, giant planet, and it was the first planet discovered by GPI (Macintosh et al. 2015). As opposed to β Pic *b* it is rather cold and seems to be cloudy (Rajan et al. 2017; Samland et al. 2017). Observations of the planet with integral field spectrographs on GPI and Keck shown in Fig. 1.5, left, show distinct spectral features in various wavelength bands, which show decent accordance with a modeled spectrum that assumes a cloudy atmosphere filled with iron and silicates. Modeling these features with varying parameters for fitted spectra lets us infer the luminosity of the planet, which in turn lets us make conclusions about the probable formation process the planet ran through. The two competing models in planet formation theory are gravitational instability (GI; Boss 1997) and core accretion (CA; Lissauer 1987). With GI, clumps of matter start forming in the outer regions of a protoplanetary disk that then contracts to become a giant planet, accreting the gas around it until its orbit is cleared out. In the CA theory, small particles of dust collide and start sticking to each other, accreting more and more material like this over time to form a planet. Evaluating a companion’s luminosity as a




Figure 1.5: *Left:* 51 Eri *b* spectral data (grey points) taken with integral field spectrographs on GPI (up to 2.5 μm) and Keck (two bands indicated with green lines). The red line represents a best fit spectral energy distribution for an atmospheric model with iron and silicate clouds. Observations with JWST will provide complementary data in the 2–20 μm wavelength regime, stretching out to the far right of this plot and beyond. *Right:* Bolometric luminosity from spectra like the one to the left as an indicator of formation theory. Dotted lines are evolutionary tracks assuming a GI start, the solid lines assume CA. The data from 51 Eri *b* is consistent with both formation models, and dynamical mass estimates with for example the methods shown in Fig. 1.4 could help to explicitize the more likely formation theory for this young, Jupiter-like planet. Source: Rajan et al. (2017).

function of its age like in Fig. 1.5, right, provides a way to discern between such different pathways to planet formation. The plot shows a number of different sub-stellar companions, including the planets of HR 8799, and 51 Eri *b*, superimposed on evolutionary tracks assuming different masses, and differentiating between formation through GI (dotted lines), and CA (solid lines). We can see how 51 Eri *b*, based on the fitted iron-silicate atmosphere spectrum, is compatible with either planet formation process. Constraining the planet mass further, for example with dynamical orbit modeling like shown in Fig. 1.4, could improve our understanding of the body’s formation process.

Contrary to the numerous disk detections with HCI instruments, the examples of directly imaged exoplanets like the ones shown above amount to only a handful of the ~ 4800 confirmed planets. This gives us a fairly small sample to conduct statistical studies on, which can provide insight into the frequency of sub-stellar companions in the proximity of the solar system and shed light on their physical nature. Over the past two decades, dedicated direct imaging surveys from the ground have targeted large samples of nearby stars to conduct comprehensive searches of the solar neighborhood, aiming to address some of these topics. The two largest ongoing surveys are the SpHere INfrared survey for Exoplanets (SHINE; Desidera et al. 2021; Langlois et al. 2021) on the SPHERE instrument of the Very Large Telescope (VLT) in Chile, and the Gemini Planet Imager Exoplanet Survey (GPIES; Nielsen et al. 2019), conducted with the Gemini Planet Imager (GPI) at the Gemini South telescope, also in Chile. They are both surveying the closest stars in a sphere of an average radius of 10 parsec around the Sun, to look for planets at separations of 10-100 au from their respective host star, probe the atmospheres of giant planets and brown dwarfs, determine the orbital properties of substellar companions and investigate their formation and evolution history (Vigan, A. et al. 2021). Shown in Fig. 1.6, we can see a comparison of companion detections around sun-like stars (F, G and K stars) from the SHINE survey with two numerical models for planet formation via GI and CA. We can see that these two pathways to planet formation produce fundamentally different planet populations: GI, marked with shaded blue dots in the plot (the shade strength indicating the density of companions) tends to form massive planets at larger separations, while CA, marked with red and yellow dots, produces the opposite. The black contour lines mark the “depth of search” of the SHINE survey, meaning how many stars SPHERE is sensitive to as a function of mass and semi-major axis, and the grey points with error bars are data points from five companion detections around F, G and K in this context.

The extended blue area marking the GI population in this parameter space overlaps with the high




Figure 1.6: Comparing SPHERE detection limits in the SHINE survey to two numerical models of planet formation, GI and CA. The GI simulation results are plotted with blue, the CA results with red dots, in the space of companion mass (in M_{Jup}) vs. semi-major axis (in au). The five grey points with error bars indicate companion detections around F, G and K stars. The black contours show SPHERE detection limits, with an indication of the number of stars to which the survey is sensitive to sub-stellar companions. A potential exoEarth with a mass of $\sim 10^{-3} M_{Jup}$ at 1 au remains far away from these detection limits. Source: Vigan, A. et al. (2021).

sensitivity regions of SPHERE in this plot, and yet, not very many sub-stellar companions were detected with the SHINE survey. This suggests that the GI formation does not seem to be the prevalent planet formation process. Results like this are major information that helps us constrain what we know about the mechanisms of planet formation, a field in which many questions remain unanswered (Wagner et al. 2019). Additionally, the data in Fig. 1.6 demonstrates point-blank how Earth-like exoplanets are very much out of reach of current HCI instruments. An Earth-analogue with a mass of $10^{-3} M_{Jup}$ and a semi-major axis of 1 au drops off the plot entirely, placing itself far away from any sensible detection limits today’s HCI instruments can claim.

Expanding HCI performance to find exoEarths

It is here where it becomes obvious that even if there are planets out there that are similar to our own, we currently do not possess the technical capabilities to find them. At this point, it is worth taking a step back and remind ourselves of what exactly it is that we define as an “exoEarth”, what observational capabilities we need to be able to detect and characterize such bodies, and how this compares to our current instrumental limitations.

The only life forms we know of are completely encapsulated on a single planet, the Earth. While there might well be living beings elsewhere in the universe, we do not, in fact, have any idea about what they are like, and what conditions they require for survival. In making an educated guess, astronomers thus decided to extrapolate from what we know about the Earth and conditions here in order to define potentially habitable extrasolar planets - exoEarths. Concretely, the search for habitable planets means that we are targeting small, rocky planets in what is called the habitable zone (Kopparapu 2018), the area around a star that allows for moderate enough temperatures such that water would exist in its liquid form, like we have it on Earth where it forms the basis of all life. Not only liquid water (H_2O), but various




Figure 1.7: Our goal: observe an exoEarth with confirmed biomarkers. *Left:* Simulated image of a solar system analogue at a distance of 12.5 pc as seen with the 15 m diameter LUVOIR A space telescope, using an apodized pupil Lyot coronagraph (APLC). After a realistic application of starlight suppression with a coronagraph, WFE influence, noise and post-processing, the planets become clearly visible, among them an Earth-twin. *Right:* Simulated spectrum of an exoEarth with ECLIPS, the coronagraphic imaging instrument on LUVOIR, modeled after the Earth as we know it today. Data like these will allow us to measure abundances of molecules in the atmospheres of planets, with the goal to identify habitable conditions. Shaded regions mark 20% bandpasses in the near-UV (blue) and optical (green), and 10% bandpasses in the near-IR, with the vertical lines in the bandpass centers. Source: The LUVOIR Team (2019).

atmospheric biomarkers such as molecular oxygen (O_2) or ozone (O_3) are also considered as indicators for possible life. The extent and position of the habitable zone depends on the stellar type and distance to the host star, so we search for companions around solar-type stars, on orbits of 1 au. All of this brings us back to the challenges of direct imaging discussed in Sec. 1.1.2, and aptly displayed in Fig. 1.1, showing the famous flux ratio of 10^{-10} for an exoEarth in the visible wavelength range. Since a 1 au planet will have a shorter apparent separation in the sky the farther away its whole stellar system is located from our own solar system, it is easier to achieve the optical performance for such detections on stars in the solar neighborhood, at around 10 pc, where the apparent distance remains comparatively large, about 1 arcsec. The scientific motivation could not be simpler - take our own world and look for an equivalent one in our direct vicinity - but the technical implications pose significant instrumental challenges on our observatories. By overcoming these challenges, what we hope for is to one day detect an image like shown in Fig. 1.7, left, and to obtain spectra allowing us to detect biomarker molecules and potentially habitable conditions.

A direct instinct is to build larger telescopes, and design more performant HCI instruments, with coronagraphs that achieve better contrast at smaller inner working angles (IWA), WFS&C systems that can accurately null the light in the DH, and post-processing techniques that squeeze out the last bit of information from our observations. The extremely large telescopes (ELTs) on the ground, planned to see first light in the next decades pursue exactly this goal. With primary mirrors of 30–40 m, they will get close to being able to detect Earth-like planets, but will be limited to the nearest M-type (cool) host stars, which are less massive and less luminous than the Sun (Guyon et al. 2012). The biggest obstacle for HCI observations from the ground is the protective envelope of our own planet: the atmosphere. Once the light from a faraway object in space reaches the Earth, it has to pass through the relatively thin layer of several dozens of kilometers of air before we can capture it with our detectors. During its passage, gradients in pressure and temperature, combined with moving air layers at different speeds in different directions cause the wavefront (WF) to get dynamically distorted, massively impeding on the quality of an image formed through a telescope. While the theoretically attainable resolution, the diffraction limit, is still given by Eq. 1.1, the real performance of the imaging instrument is now highly degraded. An *adaptive optics* (AO) system (Roddiier 1999) corrects these aberrations in real time to a certain degree, at typical rates of 1–2 kHz for high-contrast observations. It uses a wavefront sensor to sense the distortions, computes the appropriate correction shape for a deformable mirror (DM) in the optical path of the instrument and applies them to the system. Since this correction will never be perfect, the final imaging performance will be constrained by the fluctuation of AO residuals. In terms of contrast for ground-based instruments,

even with “extreme Adaptive Optics” (xAO), state of the art coronagraphs and advanced post-processing techniques, we are limited to brightness ratios of typically 10^{-5} - 10^{-7} , with the lower limit being reached extremely rarely, and exclusively under excellent observing conditions. Ultimately, the residual AO halo coming from uncorrected *speckles* will limit how well we can extinguish the light from the star. Since this halo is stronger closer to the central point source, high contrast on ground-based instruments is achievable only at larger angular separations around ~ 0.3 arcsec or more, which corresponds to a physical star-to-planet distance of around several au for nearby stars.

Ultimately, to reach the ambitious goal of detecting a habitable exoplanet around a Sun-like star, we need to put new class of telescopes into space. These will have to be large in order to provide the required angular resolution as well as gather the very few photons emitted by faint planets, and they will need to carry a suite of instruments with better coronagraphs and integrated wavefront control systems that will make them the most advanced high-contrast instruments we have used so far. While the development of such systems has been conducted with a separate focus on the individual components for a long time, it has become clear that the different parts of a high-contrast instrument need to be developed and optimized together, in order to reach the ultimate performance we need. This means that coronagraphs, usually devised as a suite of static masks and optical components, and wavefront sensing and control systems have to be created and operated as a unit. The raw contrast levels achieved with such a system will then be improved upon with post-processing methods that will finally be able to reveal planets at contrast levels of 10^{-10} .

Such large space telescopes need to have diameters on the order of 8-15 m, which poses challenges for the construction and launch of the observatory. To be able to use light-weight materials and to allow the telescope to fold into the launch vehicle, the primary mirrors will be segmented, with the most common designs being made of rings of hexagons, taking advantage of their geometrical packing factor. It was shown that the size of the telescope primary mirror has a direct influence on the expected yield of exoEarth candidates (Stark et al. 2019), which is one of the main drivers to push toward larger and larger primary mirror diameters. The mirror segmentation of course introduces optical complexity into the system, not only because the segment gaps will cause diffraction effects that need to be taken into account, but also because a sufficiently accurate co-phasing of the segments needs to be achieved in order to support the optical performance of the high-contrast instrument. Two contestants for the role of the next flagship mission are *LUVOIR* and *HabEx*. They are two powerful observatory concepts that differ in their respective implementation philosophy, but they share the capability to find habitable worlds with high performance high-contrast imaging.

The goal is clear: we need a new generation of telescopes in space, escaping the disruptive atmosphere. These observatories need to have large apertures creating photon collecting areas that provide a high sensitivity, which is best achievable with segmented primary mirrors. And we need advanced star attenuation systems yielding a contrast of 10^{-10} at angular separations as small as 0.1 arcsec in order to reach the part in the parameter space that contains potentially habitable exoEarths. The solutions catering to these needs will have to be integrated HCI instruments with state of the art coronagraphs, WFS&C systems, control architectures, observation strategies and post-processing methods, all designed and optimized as a unit to allow us to eventually find other life in the universe.

1.2 Current and future landscape of high-contrast imaging facilities

The angular resolution of an imaging instrument, and hence its capability to detect faint, close-in exoplanets depends directly on the primary diameter of the telescope. This is why, starting from telescope lenses less than half a dozen centimeters wide at the time of Galileo, the zoo of astronomical observatories has not only grown massively in numbers, but also in primary mirror diameters. Big telescopes not only improve our imaging resolution, but they also have larger collecting areas, which provide an increased sensitivity to faint signals.

Coronagraphs installed on instruments provide high-contrast capabilities on a telescope. A coronagraph is usually designed and built as part of a science instrument installed on a telescope, which can also include other analysis devices, for example spectrographs. A single telescope usually hosts more than one

science instrument, with each of them branching off part of the light that hits the telescope primary mirror into its own optical path and detector. Depending on the design and operation of the instruments, there are often several coronagraphic observing modes available on a given observatory. It is known from first principles that the geometry of the telescope entrance pupil largely impacts the image formation process and coronagraph performance. So, while each instrument has its own optical setup, control components and science objective, each of them will be equally influenced by the primary mirror geometry, its central obstruction and support struts.

Most ground-based instruments in operation today have been installed on their telescopes a long time after the observatories have been initially designed and constructed, or they have been significantly upgraded from their initial form. This means that many coronagraphs on these instruments were built under practical constraints, like accessibility of focal and pupil planes, possible alignment precision of optical masks, or the feasibility to integrate a focal-plane wavefront sensing system. However, to truly exploit the full range of high-contrast applications, whose limits are set by basic optical principles rather than engineering constraints, future telescopes need to be built from the ground up with an HCI system as part of the design. Looking ahead, high-contrast instruments are being designed as a unity (Mawet et al. 2012): the telescope structure, the static coronagraph components and active WFS&C, as well as post-processing techniques are regarded together in order to build instruments capable of reaching the lofty goal of exo-Earth detection and characterization. This is why the future large missions with high-contrast instruments we are preparing today, both on the ground as well as in space, will be opening the door to new depths in high-contrast imaging science.

1.2.1 Ground-based observatories

Telescopes on the ground have seen a long history of new technology being used to create more powerful observatories over time. The big advantage of ground-based telescopes is the access to their respective science instruments. They can be easily maintained and fixed, and after their life cycle comes to an end, they are freely replaced with or upgraded to newer generation devices. This provides a great flexibility in terms of exploring new hardware solutions, and the first-generation high-contrast imagers provided us with a leap in the imaging performance of faint objects. The first directly imaged exoplanet in 2004 (Chauvin et al. 2004) was observed with instruments on the Very Large Telescope (VLT), operated by the European Southern Observatory (ESO) in Chile. The VLT consists of four equal unit telescopes (UT), each with a 8.2 m circular primary mirror, who all saw first-light between 1998 and 2000. The NACO instrument on UT4, consisting of the Nasmyth Adaptive Optics System (NAOS) and the Near-Infrared Imager and Spectrograph (CONICA; Lenzen et al. 2003; Rousset et al. 2003), equipped with a classical Lyot coronagraph (CLC), provided the data for the 2004 detection. After moving NACO to UT1 in 2014, it was decommissioned in 2019 and is no longer in use. A direct sibling to VLT is the Gemini Observatory, consisting of two 8.1 m twin mirrors, one located in Chile (Gemini South), and one on the mountain top of Mauna Kea in Hawai'i (Gemini North). Its instruments Near-Infrared Coronagraphic Imager (NICI, decommissioned; Chun et al. 2008) on the South telescope and Near-Infrared Imager (NIRI) in the North (Hodapp et al. 2003) both used various FPM designs for coronagraphic imaging with a CLC.

The first coronagraphs to be used on ground-based systems were classical Lyot coronagraphs with adaptations of the focal-plane mask and Lyot stop to optimize them for any given observatory, telescope structure and shape, instrument, and science case. Over time, many more types of coronagraphs were deployed, especially in second-generation ground-based imagers, and today we are looking at a variation of different types and flavours of coronagraphs on many different systems. I describe those, together with the most prevalent WFS&C algorithms, in Sec. 1.3.

Current state of the art facilities

One of the most fruitfully operating high-contrast imagers today is the *Spectro-Polarimetric High-contrast Exoplanet REsearch* (SPHERE) instrument on UT3 of the VLT (Beuzit et al. 2019). It consists of three sub-systems, the Infrared dual imaging spectrograph (IRDIS; Langlois et al. 2014, 2013), the Integral Field Spectrograph (IFS; Claudi et al. 2008) and the Zurich imaging polarimeter (ZIMPOL; Thalmann et al. 2008).

Available to both of the two near-infrared instruments IRDIS and IFS, SPHERE is equipped with a whole set of masks for half-wave four-quadrant phase mask (FQPM) coronagraphs and apodized pupil Lyot coronagraphs (APLCs) (Guerra et al. 2011; Carbillet et al. 2011), while ZIMPOL has various masks available for CLCs and FQPM coronagraphs (Schmid et al. 2018). Focal plane WFS has been tested on an internal source for coronagraphic phase estimation with COFFEE (Paul et al. 2014b), and an implementation of a DH algorithm with pair-wise (PW) sensing and electric field conjugation (EFC) has been demonstrated successfully, also on the internal source (Potier et al. 2020b). After seeing first-light in 2015, an upgrade to SPHERE is being designed in order to push for even smaller inner working angles and higher spectral resolutions (Boccaletti et al. 2020a).

The dedicated exoplanet hunter on Gemini is the *Gemini planet imager* (GPI; Macintosh et al. 2018, 2014), which saw first-light in 2013. It is equipped with an APLC and is also ramping up for an upgrade to GPI 2.0 (Chilcote et al. 2020), which will include new APLC designs with smaller inner working angles and higher robustness to low-order aberrations. The refurbished GPI instrument will be moved from Gemini South to Gemini North to make place for other science instruments not dedicated to high-contrast imaging.

Rather than performing large surveys in the solar neighborhood, the *Subaru Coronagraphic Extreme AO* (SCEXAO) system (Ahn et al. 2021; Jovanovic et al. 2016; Guyon et al. 2010b) on the 8.2 m diameter, circular Subaru Telescope on Mauna Kea is working to push the limits on inner working angles with various coronagraphs and WFS&C implementations. It is designed as a laboratory testbed with a modular setup to test out new technologies, with easy switching to on-sky validations and observations (Guyon et al. 2020; Lozi et al. 2018).

A similar concept can be found on one of the two Magellan telescopes at the Las Campanas Observatory in Chile (Shethman & Johns 2003). Installed on the 6.5 m Magellan Clay telescope, *MagAO-X* is an extreme AO system targeting the visible wavelengths (0.5–1 μm), which addresses an observation window parallel to the near-IR of GPI and SPHERE (1–2.5 μm) (Males et al. 2020, 2018). The instrument follows a general Lyot architecture with exchangeable pupil- and focal-plane masks, with a vector apodizing phase plate (vAPP) coronagraph and a phase-induced amplitude apodization complex mask coronagraph (PIAACMC) used as a baseline.

All the above ~ 8 -m-class telescopes have in common that they have monolithic primary mirrors. Manufacturing, coating and successfully polishing monolithic primaries that are larger than this has proven to be challenging to impossible. One possible solution to gain in angular resolution is interferometry in the visible and IR, where the combination of signals of spatially separated mirrors is combined to produce high resolution images. The four units of the VLT have this capability, which is referred to as the VLT interferometer (VLTI; Richichi & Percheron 2005; Glindemann et al. 2003). Especially the recently commissioned GRAVITY instrument on the VLTI performs at an extremely high angular resolution at the equivalent of a 130 m telescope (GRAVITY Collaboration et al. 2017). Similar work is being conducted with the two twin-mirrors on the Large Binocular Telescope (LBT) for the LBT interferometer (LBTI; Hinz et al. 2016; Defrère et al. 2015; Hinz et al. 2014).

While interferometers allow for really large baselines since the individual mirrors can be placed very far apart, the sensitivity of such telescopes still scales with the rather sparse collecting surface. Instead, a true single-dish, large telescope can be built by stacking together individual, hexagon-shaped mirror segments to form one big mirror. One example of these segmented observatories on the ground is the Keck telescope pair, Keck I and Keck II, situated together on Mauna Kea in Hawai'i. They consist of 36 hexagonal segments each, built to a total primary mirror diameter of 10 m. Keck II features an AO system adapted to the segmented primary (Wizinowich et al. 2000) and now hosts the *Keck Planet Imager and Characterizer* (KPIC; Jovanovic et al. 2019; Mawet et al. 2016). This instrument incorporates new wavefront sensors and high-dispersion coronagraphy, and will substantially inform more advanced HCI concepts for large segmented telescopes in the future.

Future extremely large telescopes

In the next decade, we are looking at a leap in available ground-based telescope diameters and collecting areas, together with new observing technologies, introduced by the era of extremely large telescopes. The

players of this new generation of Earth-bound observatories are the *Extremely Large Telescope* (ELT²) by ESO, the *Thirty Meter Telescope* (TMT³), and the *Giant Magellan Telescope* (GMT⁴), both US-lead efforts.

With its 798 hexagonally shaped segments, composed to a 39.3 m diameter primary mirror, the ELT will be located in the southern hemisphere in Chile (Ramsay et al. 2020). With MICADO (Davies et al. 2021, 2018), HARMONI (Thatte et al. 2021) and METIS (Brandl et al. 2021), and the dedicated AO system MAORY (Ciliegi et al. 2021), it will host a whole family of first-light instruments including high performance coronagraphs for the detection and characterization of exoplanets (Houllé et al. 2021; Carlomagno et al. 2020; Baudoz et al. 2014). Further, a dedicated exoplanet hunting instrument is being planned for the second-generation of ELT instruments, the Planetary Camera and Spectrograph (PCS; Kasper et al. 2021), with an ambitious goal to look for biosignatures in the atmospheres of sub-Neptunes in the solar neighborhood.

The TMT will be composed of 492 hexagonal segments, which create a mirror of 30 m in diameter in Hawai'i (Simard et al. 2016), with the Planet Formation Imager (PFI; Kraus et al. 2016) being the equivalent to PCS. Breaking with the honeycomb-like hexagons used as segments, the GMT will be formed out of seven 8.4 m circular mirrors, arranged in a flower-like arrangement that will yield a maximum diameter of 24.5 m. Thrown into the mix of future high performance HCI instruments, the high-contrast imager TIGER (Hinz et al. 2012) will be serving its duty on this observatory in Chile.

While building our observatories on Earth enables us to support big telescopes structures like in the case of these extremely large telescopes, their optics will inevitably be impacted by the effects of our planet's atmosphere. Putting these observatories on mountain tops at high altitude, paired with xAO systems, is the best we can do to lessen its influence. However, the only solution to open the window effectively for observations of Earth-like planets is to leave the atmosphere behind entirely and put ourselves into space.

1.2.2 Space-based observatories

The first ultraviolet/optical/infrared (UVOIR) space telescope in space, revolutionizing astronomy was the *Hubble Space Telescope* (HST). It was launched and put into low-Earth orbit by a Space Shuttle in 1990, where its 2.4 m circular, monolithic primary mirror has been collecting light for its science ever since. It has been serviced by 5 shuttle missions between 1993 and 2009, during which some instruments were fixed, upgraded or swapped out for newer ones. High-contrast capabilities were provided by several of the HST instruments, like Wide Field Camera 3 (WFC3; Rajan et al. 2015), the Advanced Camera for Surveys (ACS; Krist et al. 2003), but the only currently available coronagraph in space, a simple CLC with various available masks, is installed on the Space Telescope Imaging Spectrograph (STIS; Debes et al. 2019; Grady et al. 2003). Having contributed majorly to our collective experience of putting observatories into space, and having provided the first high-contrast observations beyond the atmosphere, HST provided important learning grounds for future space missions (Krist 2004). Specifically, the next two space missions that will benefit from the lessons learned with HST will be the *James Webb Space Telescope* (JWST) and the *Roman Space Telescope* (RST).

JWST⁵ (Lightsey et al. 2012) will be the largest, and first segmented telescope that ever flew in space: its 18 hexagonal segments will form a deployable primary mirror with a diameter of 6.5 m, a mighty successor to HST. Covering IR wavelengths, JWST will have a collecting area that is six times larger, a field of view about 15 times wider, and a significantly increased spatial resolution. After launch, it will be on its way to the Lagrange 2 (L2) point between the Earth and Sun, a point in space 1.5 million km away from us, orbiting synchronously with the Earth around the Sun. Shooting a segmented mirror into space poses completely new challenges, addressed in Chapter 1.3.2. Two of its four instruments, the Near Infrared Camera (NIRCam) and the Mid-Infrared Instrument (MIRI), will include coronagraphs that will be used for the characterization of exoplanets (Perrin et al. 2018; Boccaletti et al. 2015a). While this unique observatory will address a wide range of new and exciting science, it will also be an indispensable learning opportunity

²<https://elt.eso.org/>

³<https://www.tmt.org/>

⁴<https://www.gmto.org/>

⁵<https://www.jwst.nasa.gov/>

for even larger space telescopes in the future. The design process, project management, assembly and testing alone already provided valuable lessons learned for future projects, and the same is true for the coronagraphs, both on HST and JWST (Debes et al. 2015).

Hubble has been flying for 30 years in 2020, and Webb will finish its 6-month commissioning phase in 2022, working with unprecedented high-contrast observation modes from space. However, none of these missions to date use active WFS&C within their instruments, meaning there is not a single deformable mirror dedicated to high-contrast imaging in space as of yet. This and the inherent limitations of the JWST coronagraphs prevent us from getting to a discovery space that includes close-in, rocky exoplanets. The only solution is to venture even further in the future, and conceive space telescope designs that are even larger and more complex. To this end, the NASA Astro2020 Decadal Survey evaluated several ambitious observatory designs and their recommendation was a 6 m class UVOIR telescope with high-contrast imaging capability to launch in the 2040s. This endorsement was based on the results from two design studies that defined exoEarth detection and characterization as one of their main science goals: the *Habitable Exoplanet Observatory* (HabEx⁶; Gaudi et al. 2019), and the *Large UV Optical InfraRed Surveyor* (LUVOIR⁷; The LUVOIR Team 2019).

These two observatories show different approaches to realize space telescopes that are capable of exo-Earth detection and characterization. HabEx is a 4 m monolithic telescope with an unobscured primary mirror since it is designed as an off-axis optical configuration. Its size is limited by the available launchers it would fit in to bring it to an L2 orbit around the sun. The off-axis design was chosen in favor of the vortex coronagraph (VC) in the HabEx Coronagraph (HCG; Krist et al. 2019; Riggs et al. 2018; Ruane et al. 2018a). This instrument, together with the Starshade Instrument (SSI), is the dedicated device for direct imaging of exoplanets on HabEx. HCG was conceived with two almost-identical coronagraph channels that both contain an in-pupil and an out-of pupil DM each, as well as an out-of-band Zernike WFS. A polarization beam splitter always sends vertically polarized light to channel A and horizontally polarized light to channel B. Depending on the observing strategy, the channels provide several arrangements of visible and IR cameras, and integral field spectrographs.

LUVOIR features two distinct varieties for a wide range in design space, LUVOIR A and LUVOIR B. Both are segmented with a folded launch position, and designed as a serviceable spacecraft that will be sent to L2. While LUVOIR A has an on-axis, 15 m primary with 120 segments, LUVOIR B is a 55-segment off-axis construct, both of which are larger than JWST's mirror and significantly larger than HST's, as displayed in Fig. 1.8. The instrument on LUVOIR holding promise to reveal exoEarths is the *Extreme Coronagraph for Living Planetary Systems* (ECLIPS; Pueyo et al. 2019; Juanola-Parramon et al. 2019b). It has the same design for both LUVOIR A and B, barring some minor fold mirror adjustments in the more compact LUVOIR B layout. ECLIPS is split into three different channels; near-UV, visible, and near-IR. Similarly to HabEx, each of them contains one in-pupil and one out-of-pupil DM, available pupil and focal planes for coronagraphic masks and stops, and an out-of-band Zernike WFS to monitor slow, low- and mid-order WFE for the DM correction. For LUVOIR-A, there are three separate APLC masks with different IWA, OWA and instantaneous bandwidth, each with its own FPM. The narrow-angle APLC is typically used for spectroscopic characterization in the longer wavelengths associated with molecular oxygen and water (0.76 μm and 0.94 μm). Planet detection can be performed at shorter wavelengths though, in which case we can easily tell from the basic relation of wavelength to angular resolution in Eq. 1.1 that a given angular separation between planet and star corresponds to a larger IWA in resolution units of λ/D . Having a larger IWA means using a larger FPM, and this in turn relaxes the amount of concealed area in the apodizer, increasing the coronagraph throughput significantly. In this way, the three APLC designs apply to different use-cases. In the LUVOIR A setup, ECLIPS includes a vortex coronagraph: it has a better throughput at small IWA but performs poorly on resolved stars and when there is a central obscuration, which is why it will only be used for the most distant, unresolved target stars. The LUVOIR B coronagraph is a DM-assisted vortex coronagraph (DMVC) and has a higher core throughput by a factor of two compared to the APLCs; it also has a secondary coronagraph option with an APLC. Each channel has a set of different spectral filters, apodizers, focal-plane masks and Lyot stops to enable differently optimized dark holes in

⁶<https://www.jpl.nasa.gov/habex/>

⁷<https://www.luvoirtelescope.org/>




Figure 1.8: Primary mirror comparison of monolithic and segmented space telescopes. While HST and JWST rely heavily on passive stabilization, future observatories like LUVOIR will fly full-fledged AO systems for integrated wavefront control. Adapted from The LUVOIR Team (2019).

the focal plane. All of these coronagraphs are designed to reach a raw contrast of 10^{-10} .

Considering the fundamentally different observatory designs between HabEx and LUVOIR, we notice that the general blueprint of their coronagraphy instruments is fairly similar. Both will reach the magical 10^{-10} limit in contrast, and both will certainly do so assisted by active WFS&C techniques using multiple DMs, on an observatory that will deploy various strategies in order to keep its wavefront ultra-stable over the course of an observation. However, the difference between the two telescope concepts becomes more obvious if we look at their scientific return. One metric in particular is *exoplanet yield*, the predicted number of detected exoplanets as a function of relevant observational, instrumental and astrophysical parameters (e.g., mission life time, coronagraph IWA, OWA, resolution and throughput, target star list, exoplanet occurrence rates, habitable zone boundaries, etc.). It turns out that the single most influential parameter on exoplanet yield is the inscribed parameter of the primary mirror, which in turn depends on the telescope geometry (Stark et al. 2019). The relation between yield and inscribed telescope diameter is shown in Fig. 1.9: The green curves depict HabEx-like off-axis, monolithic telescopes with charge 6 VCs. The red curves describe off-axis, segmented LUVOIR B-like architectures with DMVCs, and the blue curve expresses LUVOIR A-like on-axis (obscured) segmented designs with APLCs. The spread between equally colored curves pertains to different observing scenarios of the same telescope, and the curves stop at their assumed largest feasible size for that telescope design.

The conclusions from this plot are weighty: The constant offset of the green and red curves with respect to the blue ones indicate that coronagraphs generally perform better with unobscured telescopes, while the almost seamless connection of the green and red curves means that there is almost no penalty introduced by segmentation. The big advantage of on-axis designs seems to be plainly the technical feasibility of building them much larger than their off-axis counterparts. While high-performance coronagraphs can indeed be designed for such large, on-axis observatories, there remains a debt due to their central obscuration, marked with the blue striped area. All three of the observatories discussed above are marked in this plot: HabEx, LUVOIR B and LUVOIR A. The conclusion about yield seems simple: the larger your telescope, the more planets you will find. What is not factored into this is the fact that larger and more




Figure 1.9: ExoEarth candidate yield as a function of telescope inscribed diameter. Even if we pay a price for on-axis apertures compared to off-axis designs in terms of yield, the ability to build such structures to overall larger mirrors will result in more exoEarth candidates. The three black dots indicate the HabEx and two LUVOIR architectures. Adapted from The LUVOIR Team (2019).

complex structures, for example those with segmentation, also require more complex solutions in terms of launch configuration, commissioning activities, active WFS&C, WFE stability and others. This requires a careful analysis of the impact of structure complexity on science yield and mission feasibility, and we need to develop design and mitigation strategies for such projects.

It will be essential for the realization of a flagship mission like HabEx or LUVOIR to mature the required technologies. On the way to do that, it will be crucial to develop high-fidelity models of the coronagraph and WFS&C performance that are expected. To this end, we need to develop demonstrations of these technologies both in laboratory settings as well as in-orbit to further their technology readiness levels. While testbeds provide the venues to undertake laboratory work, the Nancy Grace Roman Space Telescope will be an indispensable pathfinder for the next generation of HCI instruments in space.

RST is a 2.4 m monolithic telescope, much like HST. The main science goal of RST is cosmology and therefore it is a wide-field infrared observatory. Unfortunately, it holds a much more unfriendly aperture geometry than HST due to the very large central obscuration, held in place by a whole array of thick support struts. On board, it will carry the Coronagraph Instrument (CGI) that features a hybrid Lyot coronagraph (HLC) as its baseline coronagraph with supported modes for shaped pupil coronagraphs (SPC; Riggs et al. 2021). CGI's requirement is to demonstrate 10^{-7} contrast in space, with the goal to reach 10^{-9} (Kasdin et al. 2020; Mennesson et al. 2018). Shown in Fig. 1.10 is a comparison between the required and predicted CGI performance with some of the instruments mentioned above, together with the real or predicted flux ratio of some known planets. It becomes clear from this figure that RST will be a major step in advancing high-contrast imaging with active WFS&C en route to the envisioned 10^{-10} observatories.

1.2.3 The role of high-contrast imaging testbeds

Moving toward the large and powerful observatories of the future that will be able to reach 10^{-10} contrast supported by active WFS&C with several DMs in space necessitates the development of the required technologies. While Roman Space Telescope CGI will constitute a major milestone in getting ready for a 10^{-10} mission, the roadmap to reach that goal is marked by the need for laboratory demonstrations. The last decade has seen an immense advancement of various coronagraph models and wavefront correction techniques which are all, on a component-level, qualified to make a high-contrast instrument on one of the




Figure 1.10: HCI performance of CGI compared to existing instruments. The lines are color-coded by bandpass center wavelength and assume an integration time of ~ 1 hr unless indicated otherwise. The different CGI observing modes are bold faced. While Roman CGI is predicted to significantly push our observing limits to deeper contrast at smaller angular separation, detection of an Earth-like exoplanet analogue around a star at a distance of 10 pc (turquoise marker) will have to wait for the next generation of large space telescopes. Adapted from Kasdin et al. (2020); Bailey et al. (2018).

flagship missions. It is now crucial to examine system-level solutions that bring the individual elements together to form a complete structure, and gauge the advantages the various systems have to offer. In order to mature and test such solutions, the community harbors several efforts in the form of HCI testbeds for laboratory research and validations of space-based coronagraphy. A recent overview of space-related testbed facilities is given in Mazoyer et al. (2019).

Overall, there are several aspects these testbeds are exploring, and conditions under which they do so:

Segmented vs. monolithic apertures: While segmentation allows for larger absolute diameters and collecting areas, such telescopes introduce complexity to the optical system. This concerns in particular the coronagraph optimization which needs to mitigate the aperture's non-circular geometry, and the WFE stability issue. It is thus important to test HCI solutions both on segmented and monolithic apertures to investigate the tradeoff between the two approaches to mirror design.

On-axis vs. off-axis telescope geometries: Coronagraph performance is significantly degraded by the presence of a central obscuration in the optical system. While certain telescope designs allow for an off-axis mirror, this becomes unfeasible for primaries beyond 10 m, and a more stable on-axis configuration is typically chosen. Current designs for future flagship missions include both obscured (on-axis) as well as unobscured (off-axis) telescope designs, on which different coronagraph types perform differently well.

Coronagraph design: Between raw contrast, inner and outer working angle, robustness to aberrations, off-axis throughput, spectral bandwidth and architecture complexity, various coronagraph designs will have distinct advantages over others, especially when paired with different telescope geometries and WFS&C architectures.

WFS&C strategies: The correction of static, and in particular, dynamic wavefront errors requires sensing and control methods with DMs. The next generation of large space telescopes will all include

multiple control loops on a whole range of spatial and temporal frequencies. As a result, there will be multiple sensor sub-systems for which there needs to be appropriate communications and offloading between them. A number of WFS&C strategies exist that are viable options for active control on a large space-based telescope, and a major point in lab testing is to compare them to each other, integrate them into a multi-level sensing and control system, and understand their relationship with the other parts of an HCI instrument.

Wavelength coverage: The spectral characterization of exoplanets requires a high contrast over a large band of wavelengths. This means that the electric field in the image plane needs to be estimated and controlled in broadband light with a bandwidth of 10–20%. While coronagraphs and WFS&C techniques can be optimized fairly easily for monochromatic light in a band of only 1–3%, which is suitable for observations in very narrow wavelength filters, it is necessary to demonstrate high-contrast imaging methods over a broad wavelength range to be representative of on-sky observations. Moreover, while it is common for ground-based instruments to be coupled with a spectroscopy component like an integral field spectrograph (IFS), this yet has to be demonstrated in the context of extremely high contrast of future space-based instruments.

Environmental conditions and target contrast: Stability in the WFE demands high thermal and mechanical stability of the laboratory environment in which the testbeds are operated. Ambient-air environments bear inherent restrictions to the ultimate contrast level they can achieve, at which point only vacuum testbeds can go beyond.

There are currently a number of testbed projects in the US and Europe investigating solutions for active coronagraphy from space, all working with their own combination of the points listed above depending on their respective goals.

The first experimental demonstrations of 10^{-10} contrast levels were achieved in the High-Contrast Imaging Testbed (HCIT) facility at JPL in Pasadena, California (USA) (Trauger & Traub 2007). In a vacuum, and in bandwidths from 2%–20%, the speckle nulling technique was used to create a one-sided (180°) DH on a band-limited Lyot coronagraph, proving that these starlight attenuation levels can be achieved (Trauger et al. 2012). Currently, the HCIT facility is home to its three main testbeds, all operating in a vacuum: the *Occulting Mask Coronagraph* (OMC) testbed, the *General Purpose Coronagraph Testbed* (GPCT) and the *Decadal Survey Testbed* (DST). The DST (Patterson et al. 2019; Seo et al. 2019) has the declared goal of running through all currently competing setups and techniques for future space-based coronagraphy missions (Ruane et al. 2019). Starting with a monolithic off-axis aperture like in the HabEx case, their team demonstrated a contrast of 4×10^{-10} with a classical Lyot coronagraph. The plan is to move to a segmented off-axis aperture, working on circular (360°) DHs and in a 10% bandwidth throughout, before including a Zernike WFS, and then expand the parameter space with on-axis apertures, larger bandwidths, various WFS&C techniques and faster, more stable corrections. The GPCT is used independently from the DST timeline to test novel coronagraphs, for example PIAACMCs and VCs, with segmented on-axis apertures, and for aggressive inner working angles. The OMC testbed is used to demonstrate the capabilities of the RST/CGI coronagraphs and WFS&C loops. Both main coronagraph designs for CGI, the HLC and the SPC, have been used to reach a contrast on the order of 10^{-9} with the obscured RST pupil (Cady et al. 2017; Seo et al. 2017). With its operations, it is preparing crucial for features space-based high-contrast missions, like two-DM control and stable WFS&C.

Another testbed aiming to demonstrate and assess HCI methods for RST/CGI is the “*Très Haute Dynamique 2*” (THD2) testbed in Paris, France (Baudoz et al. 2018a,b). It has been testing various phase-mask coronagraphs at large bandwidths (Galicher et al. 2020; Patru et al. 2018; Delorme et al. 2016) and reached the deepest in-air contrast to date with 3×10^{-9} . While most other testbeds all use the pair-wise estimator to sense the electric field in the focal plane, and EFC as a controller, one of the declared goals of THD2 is to compare various WFS&C techniques against each other (Potier et al. 2020a; Herscovici-Schiller et al. 2018; Paul et al. 2014b; Mazoyer et al. 2013). The THD2 testbed is working with unsegmented on-axis apertures, with recent adaptations being made to the entrance pupil and DMs in order to make it a truthful RST demonstrator.

Two testbeds that are dedicating their work specifically to segmented aperture telescopes are the *High-Contrast Spectroscopy Testbed for Segmented Telescopes* (HCST; Llop-Sayson et al. 2020) at Caltech in Pasadena, California (USA), and the *High-contrast imager for Complex Aperture Telescopes* (HiCAT) testbed (Soummer et al. 2018) at STScI in Baltimore, Maryland (USA). The HCST is testing apodized vortex coronagraphs (AVC) on off-axis telescope designs like LUVVOIR B, and they are preparing a single mode fiber spectrograph for high-dispersion coronagraphy (Llop-Sayson et al. 2019). HiCAT is a system-level demonstrator for on-axis segmented telescopes like LUVVOIR A. Implementing two-DM control, they use an additional segmented deformable mirror with 37 segments to simulate a segmented telescope pupil, working with APLCs. Appendix B is dedicated to this testbed as work on it constituted a significant part of this thesis.

Some additional noteworthy HCI testbeds are the *Ames Coronagraph Experiment* (ACE) located at NASA Ames in California (Belikov et al. 2012), the *Santa Cruz Extreme AO Lab* (SEAL) in Santa Cruz, California (Jensen-Clem et al. 2021), and the *Segmented Pupil Experiment for Exoplanet Detection* (SPEED) in Nice, France (Beaulieu et al. 2020). ACE is putting a focus on high contrast particularly at small angles, using PIAACMC coronagraphs on clear apertures, and the development of novel WFS&C techniques. The SEAL testbed supports ground-based AO and HCI activities on segmented telescopes, like Keck or the ELT. SPEED is a segmented testbed with dedicated experiments for ground-based telescopes.

1.3 High-contrast imaging in theory and implementation

Capturing photons to create images of faint objects in the sky for scientific analysis is an involved process that includes various significant components. First, for the purpose of high-contrast imaging, the optical system will be equipped with a coronagraph, providing the best possible static solution for starlight suppression. In an ideal case without aberrations and noise, theoretical designs are well capable of reaching contrast levels beyond the required 10^{-10} for exoEarth detection and characterization, both in monochromatic as well as in 10%–15% broadband light. This means that all the light residuals due to diffraction are in theory cancelled out at a satisfactory level. In practice though, the initially aberration-free flat wavefront will be exposed to aberration sources that introduce perturbations to the light that will be transmitted by the coronagraph, resulting in a sea of speckles around the optical axis. Wavefront aberration sources can be both internal and external to the instrument itself. Manufacturing defects, polishing errors on optical surfaces as well as alignment errors contribute to the dynamic deterioration of wavefront quality, as thermal expansion and contraction will make errors introduced by mechanical components vary over time, usually on the order of minutes to hours. On top of this, ground-based instruments are heavily impacted by the Earth's atmosphere, which causes very fast-evolving wavefront errors (on the millisecond timescale) due to a changing index of refraction in different atmospheric layers, caused by temperature fluctuations. This requires active wavefront sensing and control (WFS&C) as a second major component in a high-contrast instrument, where active optical components like deformable mirrors (DMs) are used to control the wavefront as prescribed by algorithms that are informed by sensing loops. Not only can such WFS&C systems reduce the degrading impact caused by aberrations, but they can also support the static coronagraph in creating even better contrast levels with so called DH algorithms. Such methods usually work in a loop of consecutive focal-plane sensing and DM correction steps that correct for aberrations and drive down the image-plane contrast. Third, the focal plane images captured by the science camera can be enhanced with various post-processing techniques, further correcting for errors and noise from various components in our optical system.

This thesis puts a focus on the first and second components of this list, static coronagraphy and active wavefront sensing and control, and further investigates the question of wavefront stability on large space telescopes. Even after careful deployment of active wavefront control, there will be drifts occurring in the system, degrading the contrast performance of the coronagraph. These residual errors stem from thermal instabilities and mechanical warping. These can be limited by building observatories that perform at a certain stability level in terms of mechanical and thermal drifts throughout its entire structure. Deriving requirements for this drift stability is a critical aspect in the design process of large segmented telescopes. In this section, I present the basic principles of image formation, coronagraphy, WFS&C and wavefront

stability requirements, and introduce the formalism used throughout this thesis.

1.3.1 Image formation and coronagraphy

Direct image formation

Images are created by the propagation of light through a medium and optical components like lenses and masks that will modify it depending on their properties. Light can in general be described as a complex-valued electric field (E-field), $E(\mathbf{r})$, built from a complex-valued pupil function $A(\mathbf{r})$, a phase aberration term $\phi(\mathbf{r})$, and an amplitude aberration term $\alpha(\mathbf{r})$, where \mathbf{r} is the 2D vector coordinate in the pupil plane:

$$E(\mathbf{r}) = A(\mathbf{r})e^{\alpha(\mathbf{r})+i\phi(\mathbf{r})}. \quad (1.2)$$

The light emitted from an observed object is captured by a finite-size telescope mirror that injects it into an imaging instrument, where it gets focused onto a science detector. Describing the telescope pupil with the transmission-only function $P(\mathbf{r})$, which will be zero outside of the aperture, the image formation from the telescope pupil plane to the science focal plane can be formally expressed by the far-field Fraunhofer propagation operator \mathcal{F} , which can be approximated with a Fourier transform:

$$E(\mathbf{s}) = \mathcal{F}\{P(\mathbf{r})e^{\alpha(\mathbf{r})+i\phi(\mathbf{r})}\}, \quad (1.3)$$

with \mathbf{s} being the image plane coordinate. We can detect the intensity $I(\mathbf{s})$ produced by this electric field in the image plane with a camera, as expressed by a modulus squared:

$$I(\mathbf{s}) = |\mathcal{F}\{E(\mathbf{r})\}|^2 = |\mathcal{F}\{P(\mathbf{r})e^{\alpha(\mathbf{r})+i\phi(\mathbf{r})}\}|^2. \quad (1.4)$$

Assuming a completely flat wavefront coming from a point source at infinity, for example an unresolved star, entering an aberration-free imaging system, the response of the optical system to this signal will be described by the system's point spread function (PSF). Without any aberrations present in the system ($\alpha(\mathbf{r}) = 0$ and $\phi(\mathbf{r}) = 0$), the PSF of an imager will be purely defined by the diffraction from the hard edge of the entrance pupil $P(\mathbf{r})$, which in the case of a perfectly circular aperture results in an Airy pattern, including a bright central core and concentric Airy rings. Changing the shape of the pupil means changing the diffraction effect we see in the focal plane. In the case of a segmented telescope, the segment gaps and non-circular outer edge of the aperture will create a distinct PSF depending on the geometry of the aperture (Itoh et al. 2019; Troy & Chanan 2003; Yaitskova et al. 2003; Lightsey & Chrisp 2003). Any additional pupil features, like a central obscuration and support spiders, will introduce additional diffraction effects in the focal plane. Some examples of different aperture shapes of increasing complexity and their PSFs are shown in Fig. 1.11. A planet located next to a star will create the exact same focal plane response, but offset to an off-axis position away from the on-axis star. However, as we have seen above, the immense flux ratio between an exoplanet and a star will make the planet disappear in the diffracted light of the star. Especially with more complicated apertures like those including segments, central obscuration and spiders, the planet light will be lost in a sea of so-called speckles, blobs of light created by the on-axis star.

The classical Lyot coronagraph

The idea behind coronagraphy is to suppress the diffraction effects of the star in order to let the planet shine through and render it detectable with a sufficient signal-to-noise ratio (SNR). The classical Lyot coronagraph (CLC), first used by Bernard Lyot in 1930, does this by adding an opaque mask in the focal plane, forcing the bulk of the stellar light concentrated in the PSF core to diffract around the edges of this focal-plane mask (FPM). It ends up concentrated in a ring of light in the consecutive pupil plane, were another pupil mask, called the Lyot stop (LS) after its inventor, is installed in order to block a large part of the stellar light. Since the planet is located off-axis and thus enters the optical system at an angle, it misses the FPM and avoids being diffracted to the outside of the clear pupil. A schematic of a CLC can be seen in Fig. 1.12 with the geometrical rays for an on-axis star indicated in orange and those for an off-axis planet shown in blue.




Figure 1.11: Different aperture shapes and their PSFs in a perfect optical system. *Left:* A circular aperture results in a perfect Airy pattern. *Middle and right:* Any additional structures like segment gaps, a central obscuration and spiders will cause the on-axis light to be diffracted across the focal plane.

Each of the four planes in the classical Lyot coronagraph, denoted by A , B , C and D is described by a complex electric field E_A , E_B , E_C and E_D , whose form will be defined by all upstream masks and optical components. The intensity distributions in these planes are given by I_A , I_B , I_C and I_D , respectively. While we define a propagation from a pupil plane to a focal plane with the Fourier transform \mathcal{F} , the propagation from a focal plane to a pupil plane is then denoted with the inverse Fourier transform, \mathcal{F}^{-1} , to account for the correct change of variables and change of sign through a set of successive focal and pupil planes, like in a coronagraph. For a perfect optical system without aberrations and a monochromatic plane wave entering the instrument, there will be no phase term in the first plane A , and the electric field E_A is defined purely by the pupil function $P(\mathbf{r})$:

$$E_A(\mathbf{r}) = P(\mathbf{r}). \quad (1.5)$$

This electric field gets propagated to plane B , where the FPM is located. At that focal plane, the new electric field before the FPM, E_B^- , is described as the Fourier transform of the electric field in the previous plane, as in Eq. 1.3:

$$\begin{aligned} E_B^-(\mathbf{s}) &= \mathcal{F}\{E_A(\mathbf{r})\} \\ &= \mathcal{F}\{P(\mathbf{r})\}. \end{aligned} \quad (1.6)$$

The intensity that is formed in this plane is a direct image formed by the diffraction of the entrance pupil geometry, as described above. The first step in the starlight suppression of the CLC is to insert an FPM that is adapted to the telescope PSF, causing the light to be diffracted around the mask. The CLC uses an opaque hard-edge FPM, which is in the simplest case a straightforward circular disk. Its shape is described by $M'(\mathbf{s})$, with 0 inside the opaque mask (blocking the light) and 1 elsewhere (transmitting the light). We can use Babinet's principle to rewrite the applied mask as $M'(\mathbf{s}) = 1 - M(\mathbf{s})$, with $M(\mathbf{s})$ being 1 inside the mask and 0 elsewhere, since the diffraction pattern from an opaque body and its transparent complement are equal. This means that the electric field in the focal plane after applying the FPM, E_B^+ , is

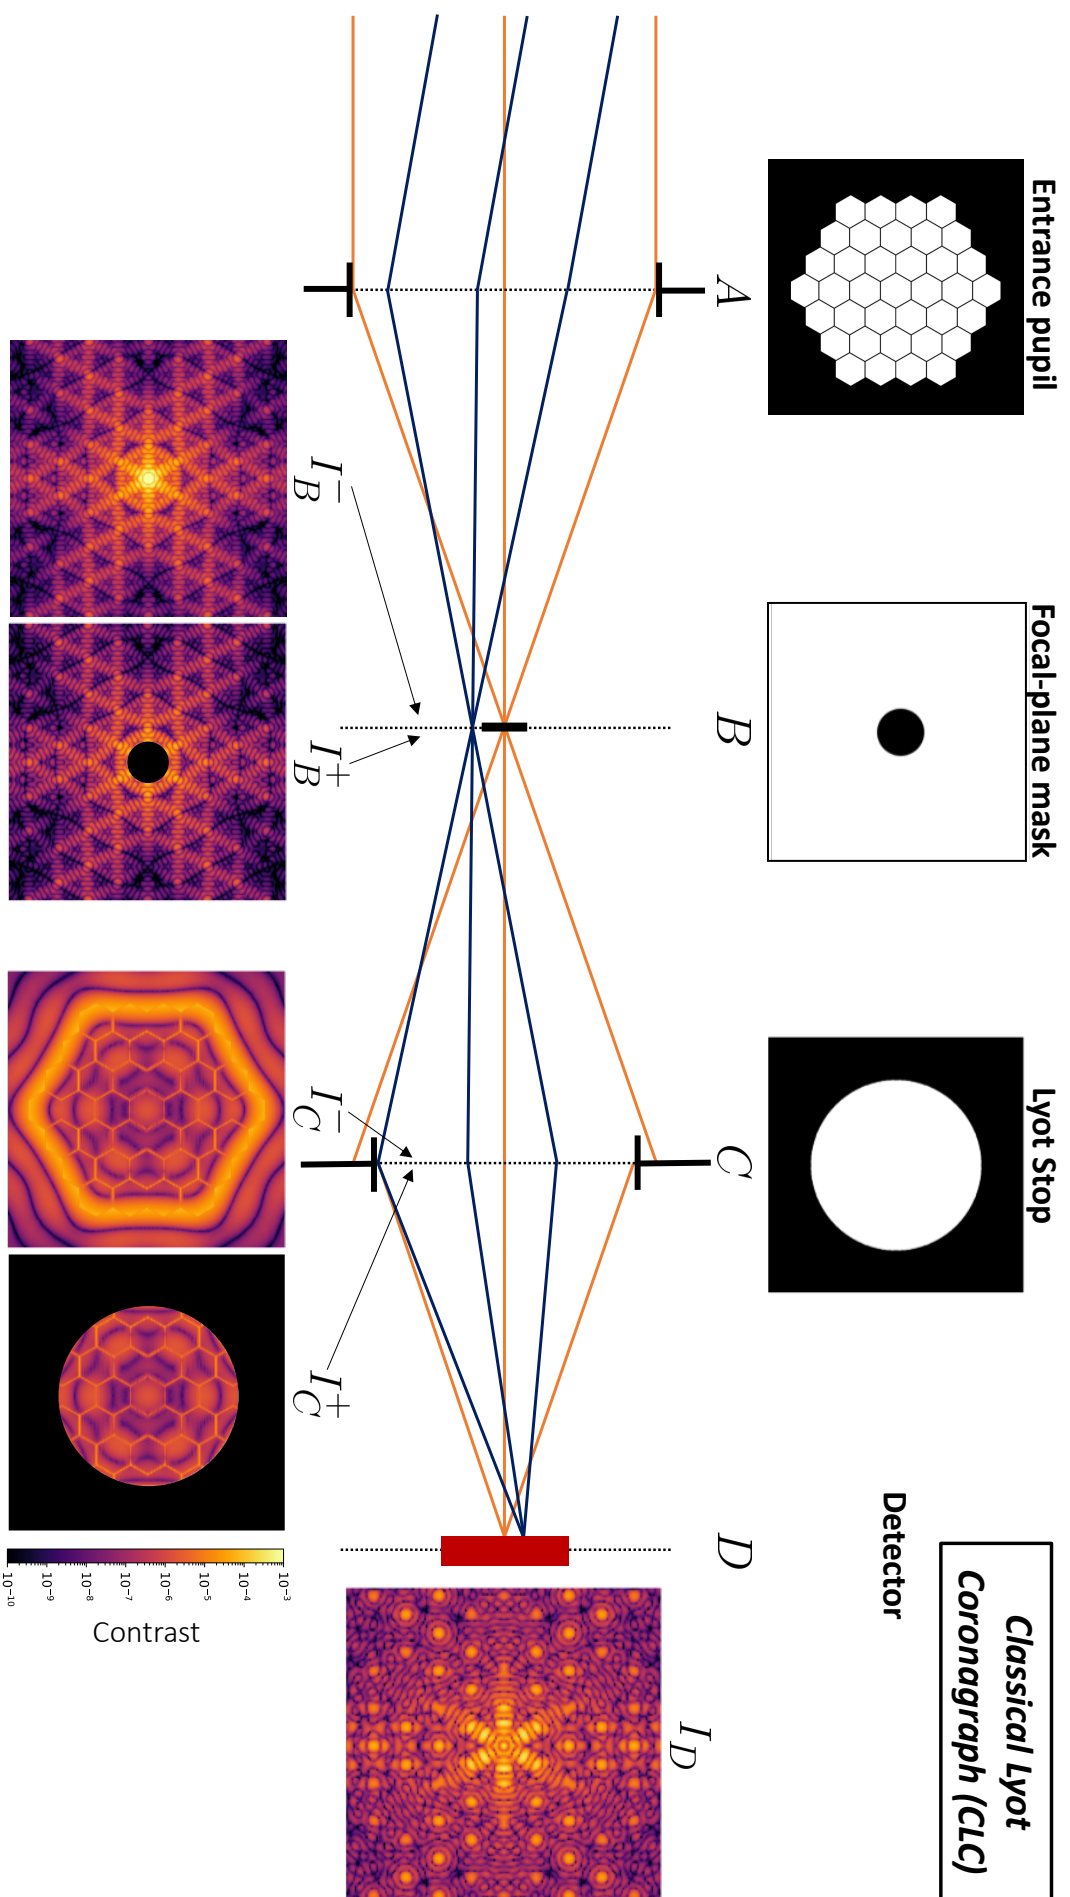


Figure 1.12: Optical layout of a classical Lyot coronagraph, with the four planes A, B, C and D as described in the text. The pupil planes A and C are assumed to have optical power similar to a converging lens. The orange rays describe the path of the on-axis starlight through the system, the blue rays denote an off-axis light source like that from an exoplanet, orders of magnitude fainter than the star. The binary masks at the top show the segmented pupil, FPM and LS that constitute the CLC. The image pairs at the bottom display the intensity before and after the FPM and LS, respectively. The image to the far right is the coronagraphic intensity on the detector. All images are plotted on the same scale and only show the on-axis light, without an off-axis source.

the multiplication of E_B^- with the coronagraphic mask:

$$\begin{aligned} E_B^+(\mathbf{s}) &= E_B^-(\mathbf{r}) \cdot M'(\mathbf{s}) \\ &= \mathcal{F}\{P(\mathbf{r})\} \cdot (1 - M(\mathbf{s})) \\ &= \mathcal{F}\{P(\mathbf{r})\} - \mathcal{F}\{P(\mathbf{r})\} \cdot M(\mathbf{s}). \end{aligned} \quad (1.7)$$

The FPM causes the on-axis stellar light to be diffracted around its edges and to interfere with the reimaged pupil in plane C , while the off-axis planet light remains almost unaffected. Before inserting the Lyot stop in plane C , the electric field E_C^- can be written as:

$$\begin{aligned} E_C^-(\mathbf{r}) &= \mathcal{F}^{-1}\{E_B^+(\mathbf{s})\} \\ &= P(\mathbf{r}) - P(\mathbf{r}) * \mathcal{F}^{-1}\{M(\mathbf{s})\} \end{aligned} \quad (1.8)$$

Effectively, plane C sees a subtraction of the diffracted wave caused by the FPM from the entrance pupil wavefront, interfering destructively. Most of the stellar light is diffracted into a ring around the edges of the pupil, which is removed by the application of a Lyot stop $L(\mathbf{r})$ in plane C , generating the electric field E_C^+ :

$$\begin{aligned} E_C^+(\mathbf{r}) &= E_C^-(\mathbf{r}) \cdot L(\mathbf{r}) \\ &= (P(\mathbf{r}) - P(\mathbf{r}) * \mathcal{F}^{-1}\{M(\mathbf{s})\}) \cdot L(\mathbf{r}). \end{aligned} \quad (1.9)$$

Finally, the coronagraphic electric field $E_D(\mathbf{s})$ is created by one last Fourier transform from the Lyot plane C to the image plane D :

$$E_D(\mathbf{s}) = \mathcal{F}\{E_C^+(\mathbf{r})\} \quad (1.10)$$

The entire coronagraph propagation from the entrance pupil in plane A to the coronagraphic plane D can be summarized with the linear operator C :

$$E_D(\mathbf{s}) = C\{E_A(\mathbf{r})\}. \quad (1.11)$$

A good coronagraph will minimize the energy in the Lyot plane (plane C), corresponding to an effective suppression of the starlight. The best solutions can be found by optimizing the FPM and Lyot Stop shapes and sizes for the given instrument. But even so, the CLC usually manages to suppress the starlight only by a factor of 10–100. Over time, the CLC has been adapted to improve its performance in various ways, giving way to a whole family of different coronagraph designs, each one presenting different advantages and disadvantages in their performance.

The Apodized Pupil Lyot Coronagraph

In the ideal case, a total extinction of starlight with a Lyot coronagraph can be achieved by rendering the two terms in Eq. 1.8 equal to each other. However, this is fundamentally impossible with the CLC because the diffracted wave from the FPM is always smooth, while the wave in the pupil is always flat, if we assume no aberrations. One solution to overcome this problem is to smoothen the pupil such that its amplitude matches better the diffracted term. This adaptation introduces the Apodized Pupil Lyot Coronagraph (APLC) where the pupil function in Eq. 1.5 gets adapted to include an apodization term $\Phi(\mathbf{r})$:

$$E_A(\mathbf{r}) = P(\mathbf{r}) \Phi(\mathbf{r}). \quad (1.12)$$

While this pupil plane apodization minimizes the residual energy in the Lyot plane greatly compared to a CLC, the extinction of starlight is not total. Such a total extinction can be achieved by apodizing other Lyot-style coronagraphs, for example the Roddier & Roddier (R&R) coronagraph that uses a phase-mask occulter instead of an opaque FPM (Roddier & Roddier 1997). While there are no perfect theoretical solutions for it, the APLC remains powerful due to its practical properties: numerically optimized apodizers can be designed with a high robustness against low-order aberrations and they work particularly well on apertures with central obscurations.

An APLC is realized with the same instrument setup like that shown in Fig. 1.12 but with an additional grey-scale or binary mask located in plane A , introducing the pupil apodization $\Phi(\mathbf{r})$, as shown in Fig. 1.13.

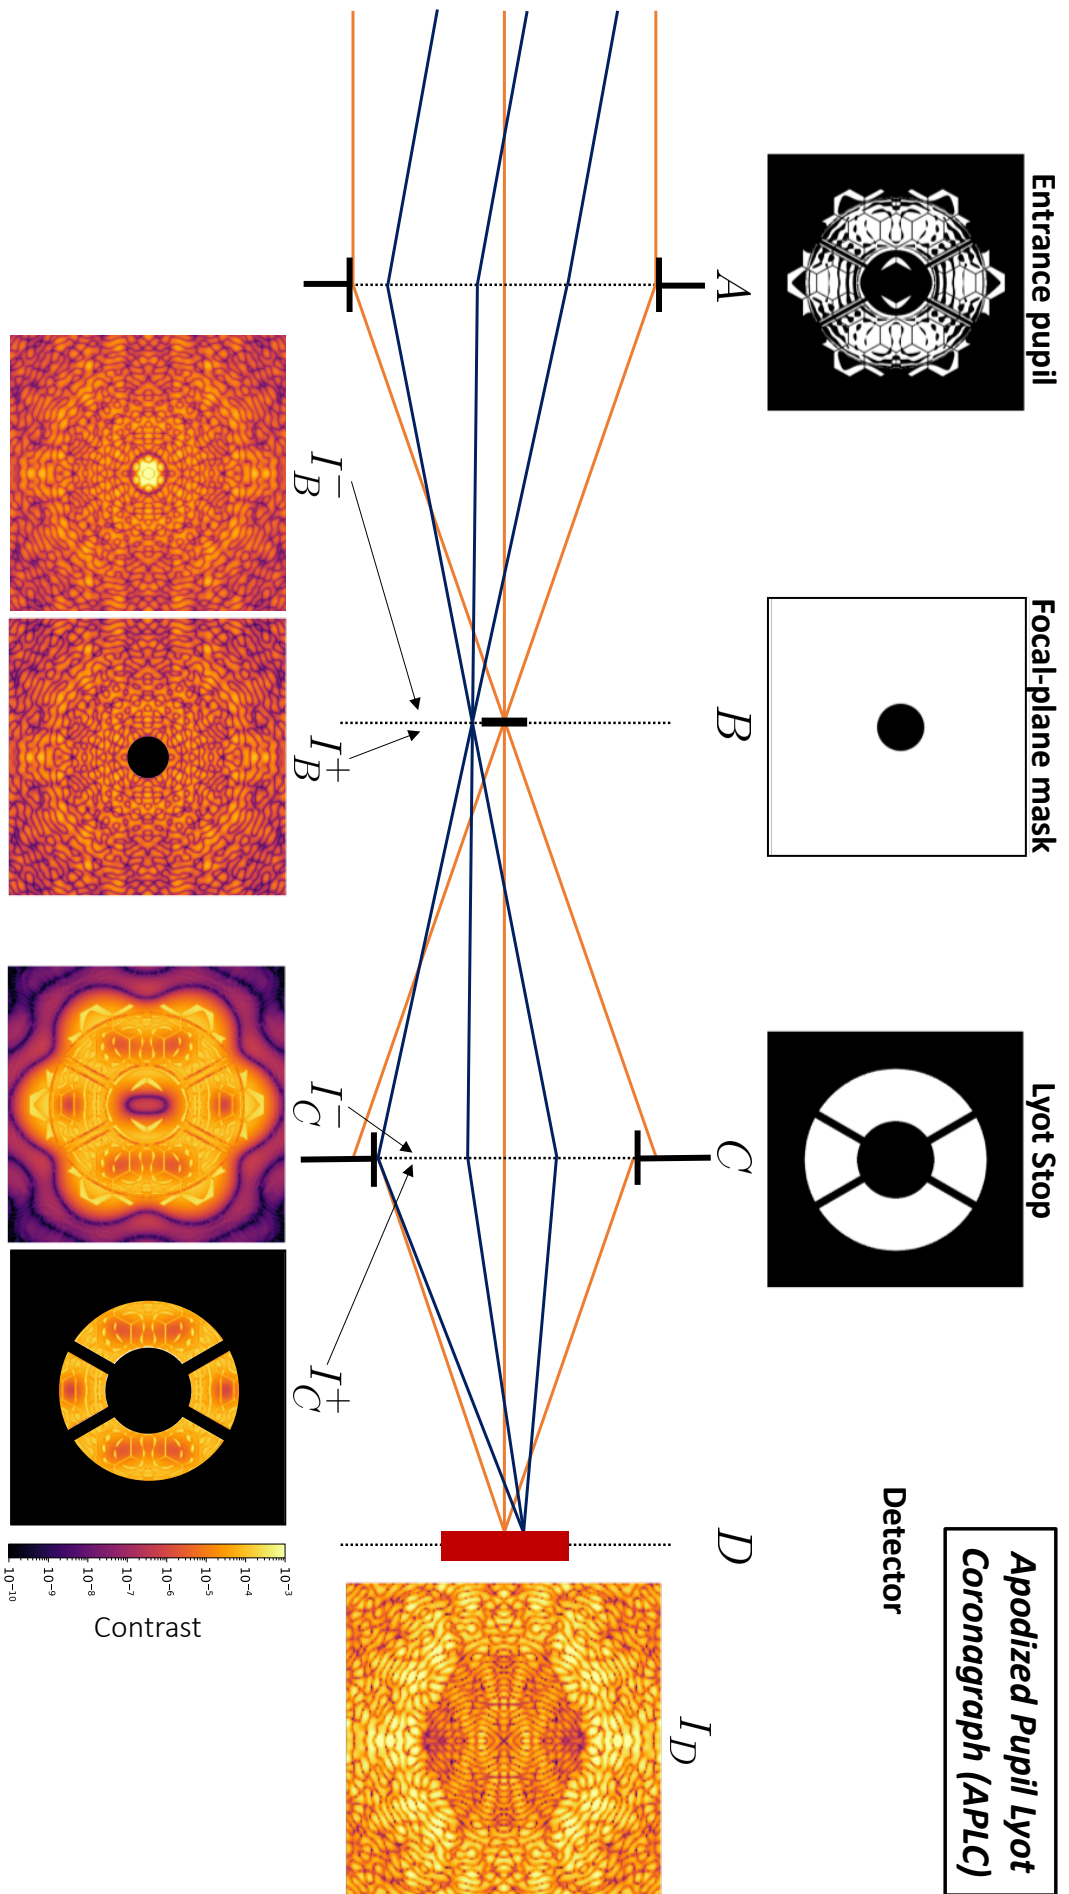


Figure 1.13: Optical layout of an apodized pupil Lyot coronagraph, with the four planes A, B, C and D equivalent to a CLC (see Fig. 1.12) but with a numerically optimized apodizer added in plane A. The pupil planes A and C are assumed to have optical power, similar to a converging lens. The orange rays describe the path of the on-axis starlight through the system, the blue rays denote an off-axis light source like that from an exoplanet, orders of magnitude fainter than the star. The binary masks at the top show the segmented pupil with the apodizer, FPM and LS that constitute the AP LC. The image pairs at the bottom display the intensity before and after the FPM and LS, respectively. The image to the far right is the coronagraphic intensity on the detector. All images are plotted on the same scale and only show the on-axis light, without an off-axis source.

The first analytical solutions have been derived for rectangular (Aime et al. 2002) and spherical telescope pupils (Soummer et al. 2003a), and they have later been generalized to arbitrary telescope geometries including a central obscuration, spiders and primary pupil segmentation (Soummer 2004). Eventually, a numerical optimization process was introduced that enables the creation of apodizers for segmented apertures that can provide contrast levels of 10^{-10} , which is sufficient for exoEarth detection and characterization around nearby stars (N'Diaye et al. 2015b, 2016). This type of optimization has been used to create APLC designs for various ground-based observatories like Gemini/GPI (Sivaramakrishnan et al. 2010) and VLT/SPHERE (Guerri et al. 2011), as well as the future space-based Large UltraViolet Optical InfraRed Surveyor (LUVOIR) telescope (The LUVOIR Team 2019).

A versatile zoo of coronagraph designs

Adding a component for pupil apodization to the CLC to create an APLC was one of many paths in the evolution of coronagraph designs, and today we are looking at a large number of coronagraph types, each presenting differences in a range of performance metrics (Ruane et al. 2018b; Guyon et al. 2006). For an excellent review of existing coronagraph designs, see Por (2020a, Chap. 1.3.1). While the work in this thesis focuses in particular on the APLC, a brief overview of the most common coronagraph types is given below.

Lyot-style coronagraphs like the APLC, based on the CLC, all feature a focal-plane occulter. Based on this simple setup, we can identify various design approaches that try to optimize either the pupil plane, or the focal plane components of the coronagraph, all the while playing with trade-offs in the coronagraph design parameter space, between contrast performance, field of view, IWA and OWA, throughput, sensitivity to aberrations and design complexity. While a pupil plane apodizer like that used in an APLC described above can improve the contrast performance of a coronagraph, it also brings the drawback of reduced throughput, as the apodization acts on the stellar and planet light equally. A similar result with a much lesser throughput loss can be achieved by using two free-form mirrors, dubbed the phase-induced amplitude apodization (PIAA) coronagraph (Guyon 2003). This type of coronagraph, especially in its evolution as a PIAA complex mask coronagraph (PIAACMC) by adding a complex FPM, enables HCI with high throughput at small IWA (Guyon et al. 2010a; Belikov et al. 2018). An implementation with similarly good throughput uses a single phase mask installed in the pupil plane to form a phase-apodized pupil Lyot coronagraph (PAPLC), and especially when paired with a knife-edge FPM, it performs at very high contrasts below $2 \lambda/D$ (Por 2020b).

The realization of the FPM can also be optimized toward different coronagraph solutions. Roddier & Roddier (1997) replaced the opaque FPM in a CLC with a π -phase mask which induces a phase shift on the on-axis light, which enables smaller IWA, and which was made achromatic by adding an additional phase ring (Soummer et al. 2003b) around the original dot of this Roddier&Roddier (R&R) coronagraph. FPMs that use both amplitude and phase are known as hybrid Lyot coronagraphs (Moody & Trauger 2007), and this design is included in the coronagraph suite of Roman CGI. There is a separate type of coronagraphs with phase mask occulters that extend out to infinity. The first of these designs was a four-quadrant phase mask (FQPM) coronagraph (Rouan et al. 2000) that is made of four equally sized, radially extending regions alternating between a phase of 0 and π . While these coronagraphs have excellent throughput due to phase-only apodization, the discontinuities between the phase regions cause the light to scatter when hitting these edges. This problem was solved by using a continuous azimuthal phase ramp instead of individual regions, which created the vortex coronagraph (VC; Foo et al. 2005; Mawet et al. 2005).

There are also coronagraphs that have no FPM and no LS at all and attenuate the light exclusively with pupil plane apodizers, making their layout very different from Lyot-style coronagraphs. Pure amplitude apodizers were initially designed fully analytically (Kasdin et al. 2003), while global optimization methods improved them significantly (Carlotti et al. 2011). Pure phase masks for pupil plane apodizers have been designed for the apodizing phase plate (APP) coronagraphs (Otten et al. 2017; Snik et al. 2012; Codona et al. 2006).

The vast parameter space for an extensive list of different coronagraph architectures provides a lot of design freedom in the definition of a HCI instrument for any given telescope. Understanding the implications of each individual design is crucial in the optimization of the imaging instrument as a unit, where the coronagraph works in close interaction with an active WFS&C system.

1.3.2 Wavefront sensing and control

Coronagraphs are highly sensitive to wavefront aberrations, so they require a careful calibration of the optical system. Wavefront sensing and control serves the purpose to modulate the electric field in an optical instrument, usually with deformable mirrors, in order to improve its performance. We can distinguish between two principal intentions of such systems: measurement and correction of the wavefront to reduce WFE, and removing speckles in the focal plane to improve contrast. An AO system aims to do the first, while the latter is usually performed with methods referred to as “DH algorithms”. They are not strictly equivalent, as nulling the wavefront does not always lead to the best possible contrast solution on a real coronagraph.

Conceptually and operationally, the WFS&C process can be split into an estimation or sensing step and a controller that are usually looped up into a number of consecutive iterations. A wavefront sensor (WFS) detects the wavefront currently propagating through the optical system and identifies phase aberrations. The wavefront controller (WFC) then applies the opposite of these aberrations in order to correct them. The process for DH algorithms is the same, except that the full electric field needs to be sensed in order to derive the optimal control step that follows to improve the coronagraph performance.

WFS&C relies on the use of deformable mirrors, which are made either segmented or with continuous control surfaces. The DM consists of a certain number of actuators that are used to deform its surface, thus actively controlling the wavefront. Since the number of these actuators on any given DM is finite within its pupil, the actuator packing determines a maximum spatial frequency in the image plane that can be controlled. A WFS can be realized with a hardware implementations set up in a separate beam path on the instrument, for example using a Shack-Hartmann WFS (Hartmann 1900a,b; Shack et al. 1971), a pyramid wavefront sensor (Ragazzoni 1996) or a Zernike wavefront sensor (Zernike & Stratton 1934). The advantage of these sensors for sensing the atmosphere is that their response is almost linear to the phase and can be quickly processed to compensate for atmospheric aberrations. The disadvantage of such setups is that they are used in a separate optical path, and any aberrations between the pickup for the sensor and the actual science focal plane are not sensed, which introduces so-called non-common path aberrations (NCPA). This means that the controller does not obtain any information on how to correct them, and the instrument performance will be limited. In particular, the greatest limitations comes from what is called *quasi-static speckles* (Bloemhof et al. 2001; Racine et al. 1999). As their name indicates, they stand in opposition to static and fast evolving speckles. Static speckles remain stable over an observation and can be subtracted with reference images, and fast varying speckles average out over the course of an observation, resulting in the AO halo. Quasi-static speckles do evolve, thus rendering them unsuitable for reference subtraction, but they are too slowly to average out.

A solution to the problem of NCPA and quasi-static speckles is focal plane wavefront sensing which performs estimations of the electric field directly in the final focal plane (Gonsalves 1982). A variety of different methods has been developed over the years (Jovanovic et al. 2018; Groff et al. 2016). This reduces NCPAs to a minimum, allowing more of the WFE to be removed. It is also a very effective method to obtain an estimate of the electric field in order to optimize the contrast in the coronagraphic focal plane. Moreover, these sensing techniques require an adaption of the science observing strategy, since some of the time will need to be allocated to WFS&C that is being performed on the same detector like the data acquisition.

In the following sections, I focus on focal plane wavefront sensing, and control, for the purpose of high-contrast imaging. Malbet et al. (1995) first mention a “dark hole” as a region in the focal plane that can be controlled and corrected by deformable mirrors, recognizing that the outer extent of this region, the outer working angle (OWA), is set by the available actuators on the DM. Building on this, a number of WFS&C algorithms have been developed that iterate between estimating the electric field and decreasing the focal plane intensity with a set of calculated DM commands, derived from the estimation (Borde & Traub 2006). Focal plane sensing can be done via spatial modulation of the E-field like with the *self-coherent camera* (SCC; Galicher et al. 2008). An alternative is temporal modulation of the E-field like with *COronagraphic Focal-plane wave-Front Estimation for Exoplanet detection* (COFFEE; Sauvage et al. 2012; Paul et al. 2013, 2014a) or the *pair-wise estimator* (PW; Give’on et al. 2011). While COFFEE relies on classic phase diversity to sense the phase through a coronagraph, PW measures the focal-plane response of well-defined DM

commands in order to break the sign degeneracy of the phase in the focal-plane E-field. On the controller side, the best established methods used today are electric field conjugation (EFC; Give'on et al. 2007) and stroke minimization (SM; Pueyo et al. 2009). While both of them aim to improve the DH contrast in each iteration step, the scalar cost function for EFC calculates the necessary DM commands by minimizing the overall energy in the DH, while SM also imposes the objective to minimize the overall stroke on the DMs in the optical system. Both the EFC and SM controllers, as well as the PW estimator are model-based, which means that their implementation requires an accurate representation of the optical model with a Jacobian matrix. The Jacobian defines the influence between all DM actuators and each image plane pixel, and the performance of said algorithms relies heavily on the exactness of the model used. Recent work has led to control algorithms that use gradient-based optimization to iteratively minimize the WF control cost function (Will et al. 2021a), which eliminates the need for a Jacobian matrix.

The work in this thesis relies on the WFS&C implementation of the PW estimator and SM controller, both of which are introduced formally in the following sections.

Electric field forward model

Before presenting the formalism of one particular estimator-controller combination, we need to establish a formal definition of the electric field in the optical system. The notation here follows the one in Groff et al. (2016). We distinguish the electric field in the pupil plane, $E_{pup,k}(\mathbf{r})$, at the WFS&C iteration k , and the electric field in the image plane, $E_{im,k}(\mathbf{s})$. We first establish our forward model by defining the different electric field contributions. The transmissive pupil function is designated with $A(\mathbf{r})$, and its amplitude and phase aberrations with $g(\mathbf{r}) = \alpha + i\beta$. The phase introduced by the DM, $\phi_k(\mathbf{r})$, consists of its phase from the previous iteration ($k - 1$) plus a small change in the current iteration k :

$$\phi_k(\mathbf{r}) = \phi_{k-1}(\mathbf{r}) + \delta\phi_k(\mathbf{r}). \quad (1.13)$$

After performing a first order approximation both for the pupil aberrations $g(\mathbf{r})$ as well as for the DM control step $\delta\phi_k(\mathbf{r})$, we can write the electric field in the pupil as:

$$\begin{aligned} E_{pup,k}(\mathbf{r}) &= A(\mathbf{r}) e^{g(\mathbf{r})+i\phi_k(\mathbf{r})} \\ &= A(\mathbf{r}) e^{i\phi_{k-1}(\mathbf{r})} [1 + g(\mathbf{r}) + i\delta\phi_k(\mathbf{r})], \end{aligned} \quad (1.14)$$

after discarding the second order term. Now if we isolate specifically the part that comes from the k^{th} control step with the DM, we obtain:

$$E_{pup,k}(\mathbf{r}) = A(\mathbf{r}) e^{i\phi_{k-1}(\mathbf{r})} [1 + g(\mathbf{r})] + iA(\mathbf{r}) e^{i\phi_{k-1}(\mathbf{r})} \delta\phi_k(\mathbf{r}). \quad (1.15)$$

This equation describes the electric field in the initial pupil plane of a coronagraph containing both the pupil plane aberrations as well as the phase introduced by the DM.

To find an expression for the image plane electric field at iteration k , $E_{im,k}(\mathbf{s})$, propagated from the pupil plane, we use the linear coronagraph operator from Eq. 1.11:

$$\begin{aligned} E_{im,k}(\mathbf{s}) &= C\{E_{pup,k}(\mathbf{r})\} \\ &= C\{A(\mathbf{r}) e^{i\phi_{k-1}(\mathbf{r})} [1 + g(\mathbf{r})]\} + iC\{A(\mathbf{r}) e^{i\phi_{k-1}(\mathbf{r})} \delta\phi_k(\mathbf{r})\}. \end{aligned} \quad (1.16)$$

We will call the first of the two terms the aberrated field in the image plane, $E_{ab,k}(\mathbf{s})$, and the second term describes the contribution to the image electric field by the applied DM command:

$$E_{im,k}(\mathbf{s}) = E_{ab,k}(\mathbf{s}) + iC\{A(\mathbf{r}) e^{i\phi_{k-1}(\mathbf{r})} \delta\phi_k(\mathbf{r})\}. \quad (1.17)$$

With a continuous face-sheet DM, its phase will be defined by the influence functions $f_q(\mathbf{r})$ of all actuators N_{act} indexed by q , with an amplitude of a_q in units of WFE. This gives us an expression for the DM phase as:

$$\phi(\mathbf{r}) = \frac{2\pi}{\lambda} \sum_{q=1}^{N_{act}} a_q f_q(\mathbf{r}). \quad (1.18)$$

The a_q are the linear actuator amplitudes at step k , so we can write each amplitude a_q at step k as the amplitude from the previous step plus a small amplitude change:

$$a_{q,k} = a_{q,k-1} + \delta a_{q,k}. \quad (1.19)$$

Equivalently, we can write the differential phase on the DM as:

$$\delta\phi = \frac{2\pi}{\lambda} \sum_{q=1}^{N_{act}} \delta a_q f_q(\mathbf{r}). \quad (1.20)$$

The full image electric field $E_{im,k}(\mathbf{s})$ can then be expressed by substituting Eq. 1.20 into Eq. 1.17:

$$E_{im,k}(\mathbf{s}) = E_{ab,k}(\mathbf{s}) + \frac{2\pi i}{\lambda} \sum_{q=1}^{N_{act}} \delta a_q C\{A(\mathbf{r}) e^{i\phi_{k-1}(\mathbf{r})} f_q(\mathbf{r})\}. \quad (1.21)$$

This equation represents the electric field in the image plane in the form of continuous functions for the fields, phases and actuator influences. Since the E-field will be detected on a camera with N_{pix} discrete pixels, indexed by p , we can put Eq. 1.21 into matrix form:

$$\mathbf{e}_{im,k} = \mathbf{e}_{ab,k} + G_k \mathbf{u}_k \quad (1.22)$$

where $\mathbf{e}_{im,k}$ is the discrete E-field vector in the image plane, $\mathbf{e}_{ab,k}$ is the E-field contribution from aberrations, both with dimensions $[p \times 1]$. The matrix G_k is the control matrix of the form $[p \times q]$ and the vector of actuator pokes \mathbf{u}_k has dimensions $[1 \times q]$. The vector \mathbf{u}_k contains all actuator pokes from 1 to N_{act} : $\mathbf{u}_k = [\delta a_1, \delta a_2, \dots, \delta a_q, \dots, \delta a_{N_{act}}]^T$. G_k is the control effect matrix/interaction matrix for the propagation of each influence function in the matrix $\mathbf{f} = [f_1, f_2, \dots, f_q, \dots, f_{N_{act}}]$, this means it tells us how the E-field changes in the image plane if we poke one single actuator in the pupil plane:

$$G_k = \frac{2\pi i}{\lambda} C\{A(\mathbf{r}) e^{i\phi_{k-1}(\mathbf{r})} \mathbf{f}\} \quad (1.23)$$

All we can measure directly in order to perform focal-plane WFS is the image plane intensity $I_{im,k}(\mathbf{s}) = |E_{im,k}(\mathbf{s})|^2$, in vectorial form expressed as $\mathbf{i}_{im,k} = |\mathbf{e}_{im,k}|^2$, with each vector component being a detector pixel. The total number of pixels N_{pix} will be contained to the high-contrast area of interest, the dark hole (DH).

Pair-wise estimation

The pair-wise estimator aims to sense the E-field from multiple discrete and noisy focal plane intensity measurements i_{im} . The estimation result is then fed to the controller to compute the optimal DM command to get a darker DH in the following iteration step of the WFS&C loop. Since one single intensity measurement is not enough to distinguish between a positive and negative phase of the corresponding E-field, we need several modulated measurements in order to obtain a clean estimate. In practice, the PW estimator modulates the electric field with well-defined DM commands in the pupil plane, so-called probes. The response in the focal plane to pairs of such probes, as modeled with Eq. 1.22, is then measured in order to estimate the electric field in the image plane (Give'on et al. 2011).

For PW estimation, we chose to sequentially apply probe commands $\pm\Psi_j(\mathbf{r})$ to the DM, where the positive and negative amplitude lead to one intensity measurement each. We call the sequential application of a positive and a negative probe a probe pair, where j indicates the number of the pair (not the individual probe). A probe can be expressed as a per-actuator DM command vector $\boldsymbol{\psi}_j$. Within a pair, probes differ only by their sign:

$$\boldsymbol{\psi}_j^\pm = \pm\boldsymbol{\psi}_j, \quad (1.24)$$

where $\boldsymbol{\psi}_j^+ = \boldsymbol{\psi}_j$ and $\boldsymbol{\psi}_j^- = -\boldsymbol{\psi}_j$. For each probe pair, we record one intensity image per probe, $\mathbf{i}_{im,k}^+$ and $\mathbf{i}_{im,k}^-$, which we can express with Eq. 1.22:

$$\mathbf{i}_{im,k,j}^\pm = |\mathbf{e}_{im,k,j}^\pm|^2 = |\mathbf{e}_{ab,k} + G_k \boldsymbol{\psi}_j^\pm|^2. \quad (1.25)$$

We can show separately for the positive and negative probe in each pair:

$$\begin{aligned}
\mathbf{i}_{im,k,j}^+ &= |\mathbf{e}_{im,k,j}^+|^2 = |\mathbf{e}_{ab,k} + G_k \boldsymbol{\psi}_j^+|^2 \\
&= |\mathbf{e}_{ab,k} + G_k \boldsymbol{\psi}_j|^2 \\
&= |\mathbf{e}_{ab,k}|^2 + |G_k \boldsymbol{\psi}_j|^2 + 2\Re\{\mathbf{e}_{ab,k}^* \circ G_k \boldsymbol{\psi}_j\} \\
\mathbf{i}_{im,k,j}^- &= |\mathbf{e}_{im,k,j}^-|^2 = |\mathbf{e}_{ab,k} + G_k \boldsymbol{\psi}_j^-|^2 \\
&= |\mathbf{e}_{ab,k} - G_k \boldsymbol{\psi}_j|^2 \\
&= |\mathbf{e}_{ab,k}|^2 + |G_k \boldsymbol{\psi}_j|^2 - 2\Re\{\mathbf{e}_{ab,k}^* \circ G_k \boldsymbol{\psi}_j\},
\end{aligned} \tag{1.26}$$

where \circ signifies element-wise multiplication. We can see the intensity from the positive and the negative probe within a pair only differ by the sign in front of the mixed term. If we define the difference $\Delta \mathbf{i}_{k,j} = \mathbf{i}_{k,j}^+ - \mathbf{i}_{k,j}^-$, only twice the mixed term remains:

$$\Delta \mathbf{i}_{k,j} = \mathbf{i}_{k,j}^+ - \mathbf{i}_{k,j}^- = 4 \Re\{\mathbf{e}_{ab,k}^* \circ G_k \boldsymbol{\psi}_j\}. \tag{1.27}$$

Using the identity $\Re\{A^*B\} = \Re\{A\}\Re\{B\} + \Im\{A\}\Im\{B\}$, we obtain:

$$\Delta \mathbf{i}_{k,j} = 4[\Re\{\mathbf{e}_{ab,k}\}\Re\{G_k \boldsymbol{\psi}_j\} + \Im\{\mathbf{e}_{ab,k}\}\Im\{G_k \boldsymbol{\psi}_j\}]. \tag{1.28}$$

Since the real and imaginary parts of the two terms are completely separated, we can write Eq. 1.28 in vector format:

$$\Delta \mathbf{i}_{k,j} = 4 \begin{pmatrix} \Re\{G_k \boldsymbol{\psi}_j\} & \Im\{G_k \boldsymbol{\psi}_j\} \\ \Re\{\mathbf{e}_{ab,k}\} & \Im\{\mathbf{e}_{ab,k}\} \end{pmatrix}. \tag{1.29}$$

We now want to solve this equation for the electric field vector, $\mathbf{e}_{ab,k}$, on the right hand side in order to obtain our estimation result for the aberrated E-field. To do that, we need to invert the matrix that holds the coronagraph propagation of the DM phase, called the observation matrix. Since we have used only a single probe pair, this is only a $[1 \times 2]$ matrix though, and an inversion cannot be done. If we increase the number of pairs to at least 2, this matrix becomes invertible. For exactly $N_{pairs} = 2$ ($j = 1, 2$), Eq. 1.29 is written as:

$$\begin{pmatrix} \Delta \mathbf{i}_{k,1} \\ \Delta \mathbf{i}_{k,2} \end{pmatrix} = 4 \begin{pmatrix} \Re\{G_k \boldsymbol{\psi}_1\} & \Im\{G_k \boldsymbol{\psi}_1\} \\ \Re\{G_k \boldsymbol{\psi}_2\} & \Im\{G_k \boldsymbol{\psi}_2\} \end{pmatrix} \begin{pmatrix} \Re\{\mathbf{e}_{ab,k}\} \\ \Im\{\mathbf{e}_{ab,k}\} \end{pmatrix}. \tag{1.30}$$

This system of equations can be extended to an arbitrary number of probe pairs $N_{pairs} \geq 2$. By calling the vector of differential intensity measurements $\mathbf{z}_k = [\Delta \mathbf{i}_{k,1}, \Delta \mathbf{i}_{k,1}]^T$, the vector of electric field components $\mathbf{x}_k = [\Re\{\mathbf{e}_{ab,k}\}, \Im\{\mathbf{e}_{ab,k}\}]^T$ and the observation matrix H_k , we can rewrite Eq. 1.30 as:

$$\mathbf{z}_k = H_k \mathbf{x}_k. \tag{1.31}$$

Finally, to obtain an estimate of the aberrated E-field, denoted by $\hat{\mathbf{x}}_k$, we need to take the pseudo-inverse of H_k :

$$\hat{\mathbf{x}}_k = (H_k^T H_k)^{-1} H_k^T \mathbf{z}_k. \tag{1.32}$$

Given that all high-performance HCI instruments will carry a set of DMs for control, electric field estimation with PW probing is a straightforward way to perform wavefront sensing. Some limitations apply in its efficiency, which are mostly given by the choice of probes used. The probe shapes in the pupil plane need to maximize the signal in the image plane and provide a well-conditioned observation matrix, which means that they need to modulate the intended control area, the DH. Traditionally, simple analytical functions in the focal plane have been used, for example rectangular regions, since their Fourier transforms are known and can easily be used to define the DM commands of the probes. Another option is to back-propagate the designated DH area in the focal plane into the DM pupil plane by using the Jacobian matrix G . In this way, it is ensured that the DH region is always fully modulated. Another influencing factor on the estimate accuracy is the probe brightness. While brighter probes reduce the variance of the electric field estimate, they cannot be too large, otherwise the higher order terms discarded in the Taylor expansion of the forward model become the dominant error.

Control with stroke minimization

Once we have an electric field estimate for a given iteration in our WFS&C loop, we want to use this information to calculate and apply a correction command with the DM in order to obtain a better average contrast in the DH. Knowing the forward model for the electric field vector $\mathbf{e}_{im,k}$ given in Eq. 1.22, with $\mathbf{e}_{ab,k}$ coming from the estimator, we can calculate the intensity vector in the image plane by taking the scalar product of $\mathbf{e}_{im,k}$ with itself:

$$\begin{aligned} \mathbf{i}_{im,k} &= \langle \mathbf{e}_{im,k}, \mathbf{e}_{im,k} \rangle \\ &= \mathbf{u}_k^T G_k^* G_k \mathbf{u}_k + 2 \mathbf{u}_k^T \Im \{ G_k^* \mathbf{e}_{ab,k} \} + \mathbf{e}_{ab,k}^* \mathbf{e}_{ab,k}. \end{aligned} \quad (1.33)$$

Concretely, the matrix $G_k^* G_k$ designates the additive contribution of the DM to the image plane intensity, at iteration k . The vector $G_k^* \mathbf{e}_{ab,k}$ represents the interaction of the electric fields from the DM with the aberrated field, and the vector $\mathbf{e}_{ab,k}^* \mathbf{e}_{ab,k}$ indicates the intensity contribution from the aberrated field $\mathbf{e}_{ab,k}$.

Since the general goal of the WFS&C loop is to improve the contrast in the DH, we need to define a control strategy to lower the image plane intensity after the coronagraph. The stroke minimization controller (Pueyo et al. 2009) defines a strategy to minimize the total stroke on the DM while subjecting the average contrast in the DH to the condition to become lower than a target contrast value chosen for that particular iteration:

$$\begin{aligned} \underset{\mathbf{u}_k}{\operatorname{argmin}} \quad & \mathbf{u}_k^T \mathbf{u}_k \\ \text{subject to} \quad & I_{DH,k} \leq c_{target,k}. \end{aligned} \quad (1.34)$$

To be able to create a cost function from this strategy, we treat the contrast condition as an equality:

$$I_{DH,k} - c_{target,k} = 0, \quad (1.35)$$

which allows us to define the quadratic cost function J_k using a Lagrange multiplier μ :

$$J_k = \mathbf{u}_k^T \mathbf{u}_k + \mu_k (I_{DH,k} - c_{target,k}). \quad (1.36)$$

To find the minimum of this cost function, we take its partial derivative $\partial J_k / \partial \mathbf{u}_k^T$ and solve for the control command \mathbf{u}_k after setting the derivative to zero:

$$\mathbf{u}_k = - \left(\frac{1}{\mu_k} \mathcal{I} + G_k^* G_k \right)^{-1} \Im \{ G_k^* \mathbf{e}_{ab,k} \}, \quad (1.37)$$

where \mathcal{I} is the identity matrix.

The controller uses this DM command solution and with initial Lagrange multiplier μ_k^0 to calculate the predicted DH intensity resulting from it. If the resulting $I_{DH,k}$ lies below the target contrast $c_{target,k}$ we demanded in this iteration, then a line search on μ_k is performed with a step size α until a solution is found that indeed brings the image plane intensity to or below the target contrast value:

$$\mu_k^n = \alpha \mu_k^{n-1}. \quad (1.38)$$

Since the measured average contrast in the DH decreases with each iteration of the WFS&C loop, the requested target contrast is usually defined relative to the current measured contrast value. In early iterations, when the residual WFE is still larger and the speckles brighter, a larger relative step size will make the contrast loop converge faster, while smaller steps later on will keep the loop stable.

1.3.3 Post-processing

The coronagraph and active WFS&C on a telescope aim to optically separate the starlight from the planet light before its registration on the detector. The resulting contrast performance, commonly referred to as the *raw contrast*, is highly limited by the slowly evolving quasi-static speckles, which are easily mistaken for an astronomical point source especially since they appear on the same angular scale. While an optimization of the starlight suppression instrument, including static masks and the WFS&C strategy, is




Figure 1.14: Coronagraphs are extremely sensitive to WFE, as shown here with a simulation of a segmented CLC on HiCAT. On the left, we can see the PSF from a noise-less, aberration-free coronagraphic image. On the right, the same setup now also contains a WFE of 15 nm rms composed of equal parts of astigmatism, coma, trefoil and a random phase screen with a power spectral density with a slope of 3. Both images are shown on the same scale.

needed to minimize speckles and reduce the detection noise, some improvement can be brought about with the use of post-processing algorithms.

To be able to deploy methods for the separation of the star signal from that of a planet or a disk, we exploit the fact that the stellar light exhibits different properties from the light of the source of interest. In this context, “reference subtraction” has a long history, with and without coronagraphs, where the goal is to identify the instrumental PSF in the science camera image so that it can be subtracted off the star to leave the planet’s signal. These methods improved significantly with the exploitation of specific diversities between the star and planet light, for example with Angular Differential Imaging (ADI; Marois et al. 2006) where the planet seemingly moves around the star from the perspective of the telescope as the sky rotates. This technique was advanced by using a maximum likelihood estimation to discriminate planet signals from speckles, dubbed ANDROMEDA (Cantalloube et al. 2015; Mugnier et al. 2009). There is a range of approaches to diversity imaging: Spectral Differential Imaging (SDI; Smith 1987; Sauvage et al. 2006) uses spectral differences between the star and planet, and Polarization Differential Imaging (PDI; Kuhn et al. 2001) their different polarization state, which is mostly used on disks. Coherent Differential Imaging (CDI; Guyon 2004) leverages interference between the starlight and the stellar core, which is not possible for the planet, and Orbital Differential Imaging (ODI; Males et al. 2015) uses the actual orbital motion of the planet around the star to extract its signal.

Making use of several diversities can be combined to enhance results (e.g., Rameau et al. 2015; Marois et al. 2010a), and the algorithms can be expanded to optimize entire libraries of reference images for optimal subtraction (Amara & Quanz 2012; Soummer et al. 2012; Pueyo et al. 2012).

1.4 Application of HCI to the case of large segmented telescopes

The PSF in Fig. 1.11, middle, illustrates the diffraction effect of a segmented aperture, producing a well defined intensity distribution in the image plane. This example, however, assumes a perfectly aligned system without any wavefront errors, which is an idealized scenario. In a real optical system, aberrations will lead to a degradation of image quality. Coronagraphs are particularly sensitive to WFE in the optical system, as they will contaminate the focal-plane image as illustrated on simulations of a segmented CLC on HiCAT in Fig. 1.14. On a telescope, there are many sources of WFE, but a big component on segmented telescopes are misalignments between the segments, a problem that monolithic telescopes do not have. Before we can even concern ourselves with their fine alignment and focal plane wavefront sensing as described in the previous section, the first stepping stone toward using a segmented telescope is to get the segments phased well enough with respect to each other. This will produce a PSF whose residual WFE is

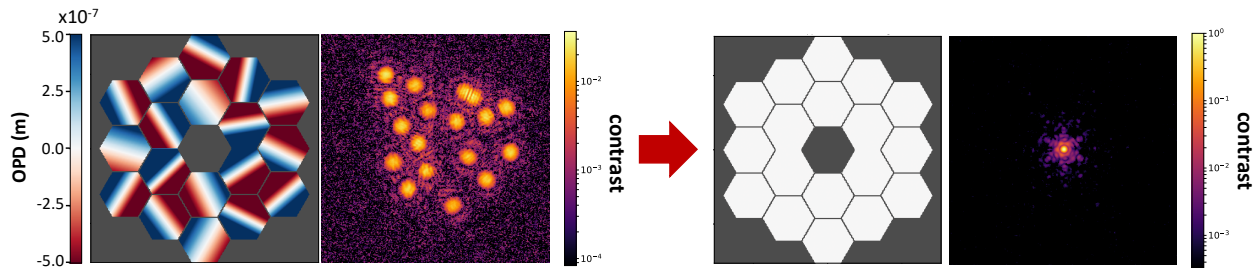


Figure 1.15: Segmented apertures require careful phasing of the individual segments in order to go from an unphased state like that shown on the left to a well-aligned telescope like that shown on the right. The WFE maps shown here indicate the segmented DM command used to produce the displayed PSFs, which still contain residual WFE. I acquired this data on the JOST testbed in the very beginnings of my PhD while investigating alignment procedures for segmented aperture telescopes (see Appendix A for more details).

within the capture range of more sensitive fine-phasing and focal plane WFS&C algorithms that optimize the coronagraphic performance.

Even after the initial telescope coarse and fine alignment is done, it is not a perfectly static structure that will be able to maintain its imaging performance. In particular the coronagraph contrast will depend on continuously executed WFS&C activities. In the following section, I briefly introduce the concepts of segment cophasing, and the significance of closed-loop WFS&C on large, segmented observatories.

1.4.1 Wavefront errors and cophasing of segments

When segmented telescopes see first light with a primary alignment, the individual segments are not cophased very well. Instead, they show inter-segment misalignments that can amount to surface errors several times the operating wavelength. The first goal during the commissioning of a segmented observatory is thus the cophasing of the segments to get as close as possible to the diffraction limit. The two states of the segmented aperture in Fig. 1.15 illustrate this: starting from completely unphased segments that produce separate “sub-PSFs”, one from each segment as shown on the left, a stacking and phasing process aims to produce a PSF with minimal residual errors, as shown on the right. There is a significant legacy for this from the Keck telescopes in Hawai‘i, where decades of work on segment phasing and segmented WFS&C strategies have been investigated and used, while continuing to look for alternatives and higher-performance methods (Chanan et al. 2000). While the segment coarse alignment is done initially and can usually be repeated as needed after a couple of months or years, WFS&C maintenance activities for fine-alignment are repeated on timescales from hours to days, depending on the telescope and its instruments.

On JWST for example, the first segmented space telescope, and the first space-based observatory to use active control components for alignment and WFS&C (Knight et al. 2012b), the surface misalignments between segments after launch and deployment are expected to be on the order of several mm, which corresponds to very large WFE at a wavelength of $2\ \mu\text{m}$. During the first months of commissioning (Perrin et al. 2016b), an extended procedure of complex steps is taken to achieve the required WFE (Acton et al. 2012). After an initial focus-sweep with the secondary mirror in order to bring the telescope to a reasonable focus, the segments need to be stacked and coarsely phased before moving on to the fine-phasing stage. Since the segments are expected to be so significantly misaligned that their light does not even overlap into a single point on the detectors (see Fig. 1.15, left), the next step will be to identify which of the “sub-PSFs” belong to which of the segments in the telescope pupil (“segment ID”). The process continues with eliminating segment-level wavefront errors before stacking them all up into a single point in the center of the field. Large piston errors between the segments are reduced with dispersed fringe sensing (Shi et al. 2004), and field-dependent errors are corrected by adjustments with the secondary mirror. Finally, the fine phasing of the telescope is performed with a phase retrieval algorithm (Dean et al. 2006), with field-dependent errors across the science instruments being addressed with a multi-instrument multi-field algorithm (Acton & Knight 2012). The expectation after completion of this lengthy process on JWST is a

fine-aligned PSF with a residual WFE on the order of 80 nm rms, similar to Fig. 1.15, right.

A bulk of the JWST commissioning process has been demonstrated on the one-meter Testbed Telescope (TBT) at Ball Aerospace (Acton et al. 2007). Such testing and validation in a laboratory environment is a crucial aspect in the development and demonstration of new algorithms. While each observatory will usually have a dedicated high-fidelity facility for simulation and experimentation, like the TBT for Webb, there are numerous non-mission-critical facilities that offer themselves for the investigation of novel methods and algorithms in WFS&C. One such testbed is the JWST Optical Simulation Testbed (JOST) at the Space Telescope Science Institute (STScI) in Baltimore, Maryland (USA) which is being used to test new WFS&C techniques for segmented telescope co-phasing (Perrin et al. 2014b; Choquet et al. 2014). Designed as a refractive analogue to a three-mirror-anastigmat, it contains a segmented DM with segments controllable in piston, tip and tilt and an adjustable lens that acts as a secondary mirror surrogate in order to investigate WFS&C techniques on a wide field of view. It currently runs a linearized algorithm for phase diversity (LAPD; Mocœur et al. 2009) that was used to align the optical components on the testbed (Egron et al. 2016, 2017a,b). Investigations about LAPD with a segmented DM on a fully aligned JOST testbed, and on a wide field of view, build Appendix A of this thesis.

Even the best practically achievable optical alignment and WFS&C strategy will leave residual errors in the imaging system that stem from drifts in various observatory components. Thermal instabilities in the telescope can cause misalignments of the secondary mirror with respect to the optical alignment which introduces global aberration modes in the pupil. Warping of the backplane structure can lead to co-phasing errors of the primary segments, adding WFE terms that degrade the performance of the imaging instrument. To be able to reach and keep the contrast levels required for exoEarth observations over the timescales of an observation, telescopes need to be designed from the beginning in a way that ensures sufficient WFE stability. In the following section, I introduce the framework and goals of designing extremely stable space telescopes.

1.4.2 Ultra-stable telescopes

Like any other mechanical structure, especially one that is exposed to the harsh environment of space, an observatory is not a perfectly stable assembly. Thermal structure deformation, mechanical vibrations and rigid-body misalignments, among others, will be constantly adding to changes in the construction that translate directly into WFEs in the optical system, on top of any static aberrations coming from surface figure errors for example. The resulting small errors can imprint features on the final image that look indistinguishable from faint exoEarth candidates. In order to minimize the number of spurious detections associated with these features, their variations need to be smaller than the challenging 10^{-10} contrast requirement. This in turn implies minimizing wavefront variations over time. Regardless of aperture (monolithic or segmented), the key challenge of any exoEarth characterization telescope will thus be to eliminate or control systematic errors across the telescope's pupil to enable reliable imaging of faint objects at contrast variations of 10^{-10} – 10^{-11} . Recent advances in coronagraphic design and adaptive optical control of systematic errors has made this a tractable engineering problem, even if controlling a large, segmented 15 m structure might seem daunting. A significant number of these innovations have already been applied to 8 m class ground-based telescopes, albeit at more modest contrasts.

Considering that even 1 pm of WFE is enough to create planet-like artifacts in the image plane, space telescopes that aim to reach exoEarth detection levels will need to be held “ultra” stable compared to the easier to do, relatively speaking, nanometer-level controlled observatories astronomers use today. There is the option to make a telescope passively stable enough so that this drift does not exceed the needed contrast level, as is done for Hubble, and even for JWST, considering its WFS&C activities are performed in *open-loop* control from the ground. However, bearing in mind that a LUVOIR-like observatory would be significantly larger than its predecessors, as we can see in Fig. 1.16, the problem seems to become exceedingly hard to solve. Indeed, a multi-meter scaled-up structure of a static, rigid telescope would require picometer-level stability over a very large aperture, and that over the course of hours, which seems discouraging.

To solve this issue, a general rethinking of the way we build and control such telescopes needs to take place. Recent mission concepts have identified “on-the-fly” WFS&C as a key aspect of the coronagraph instrument and mention the need for wavefront sensors within the exoplanet imaging instrument to op-




Figure 1.16: Comparison of observatory sizes between the Hubble Space Telescope (left), James Webb Space Telescope (middle) and LUVOIR A (right). The monolithic Hubble and segmented Webb use passive control, while LUVOIR will require active control strategies to minimize WFE drifts, and hence contrast drift. Source: The LUVOIR Team (2019).

erate in *closed-loop*. The telescope mirror cannot be considered apart or distinct from the overall optical control system of the telescope when assessing whether or not it meets science requirements: in a telescope system with wavefront-controlled adaptive optics, the static and dynamic behavior of the individual telescope components, including the mirror, all reside within the control loop. In the future, space-based telescopes will need to implement full-fledged adaptive optics systems like they have been used on the ground for decades. While those are mostly concerned with correcting aberrations caused by Earth’s atmosphere, they also correct for distortions coming from within the telescope itself. High spatial frequency distortions from the mirror support structures, and in particular dynamic misalignment of the primary segments in the case of Keck (Lamb et al. 2021), are examples of wavefront deformations arising within the telescope that their respective AO system successfully corrects for, which is crucial for coronagraphic HCI observations.

While AO control loops are fully embraced by the ground-based community, and technologists and researchers focusing on high-contrast instrument development have started adopting this more sophisticated view of controlling optical alignment (Potier et al. 2021), these methods are only slowly finding their place in the operations of space telescopes (Pueyo et al. 2021). Apart from the fact that AO relaxes the stability requirements by an order of magnitude, recent studies, notably the ones presented in this thesis, have also shown that not all segments on a segmented telescope need to be constrained equally, which offers additional relief for WFE tolerancing. Fully harnessing and applying these powerful techniques to space-based observatories will allow us to develop missions with sufficient aperture size for a significant scientific yield.

1.5 This thesis

A major part of this thesis is concerned with the definition of a general, theoretical framework for the analytical treatment of the mid-spatial-frequency WFE components associated with primary mirror segments in the stability analysis of a large, segmented telescope. It is explored how to make a differential allocation of WFE across a segmented pupil rather than following a purely global approach, supported by testbed experiments. This work is then folded into the larger picture of the treatment of ultra-stable space telescopes, and how to prepare for full adaptive optics systems in space.

Chapter 2: Analytical wavefront error tolerancing on a large segmented telescope

Chapter 2 introduces the theory of segmented telescope WFE tolerancing, based on the PASTIS analytical propagation and tolerancing model initially published by Leboulleux et al. (2018b). After extending the analytical forward propagation model to use a semi-analytic propagation matrix, I formulate the statistical interpretation of the inverted PASTIS equation. Using this statistical framework, I show how to derive segment-level WFE requirements for a chosen target contrast, and demonstrate that not all segments in

the telescope pupil need to be held equally stable.

Chapter 3: Numerical tools and simulated tolerancing results

Chapter 3 first presents the software package I developed for my work on segmented mirror tolerancing. It contains a set of optical simulators for telescopes with the capability to impose various aberration modes, including a segmented deformable mirror that can be paired with coronagraphs. The package also includes all necessary functions for the analytical tolerancing analysis.

After that, simulated results of a tolerancing analysis on JWST are shown, with a classical Lyot coronagraph on the NIRC*am* instrument. I draw a rough comparison to the real WFE expectations of JWST. The last part of this chapter presents tolerancing results for the segmented mirror on HiCAT, as performed on the HiCAT simulator. As it acts as a proper testbed emulator, I used this work to prepare the experimental validations that followed after.

Chapter 4: Experimental validations of segmented WFE tolerances on HiCAT

Chapter 4 shows PASTIS tolerancing experiments performed on the HiCAT testbed. Since the segmented deformable mirror is not the only source of WFE drifts in a laboratory environment, I provide a reformulation of the formalism from Chap. 2 to isolate the contrast contribution from the segmented DM alone. I successfully measure an experimental contrast sensitivity matrix that is used for the further tolerancing analysis. Through the model inversion, I calculate individual segment tolerances for a target contrast that is appropriate for the current contrast level on HiCAT, and validate them with Monte Carlo experiments. I then proceed to measure the coronagraphic sensitivity to all segments under a range of misalignments of the Lyot stop, and predict WFE requirements for the segmented DM for contrast levels beyond the current performance limit of the testbed.

Chapter 5: Analytical WFE tolerancing in dynamic studies for ultra-stable telescopes

Chapter 5 dives into applications of the PASTIS tolerancing framework that answer issues about picometer-level stability for future segmented telescopes. Swapping segment-level Zernikes for modes obtained from thermal modeling, I show how to apply the method to physical tolerancing applications. More globally, I show how this tolerancing approach folds into a larger paradigm of tackling WFE stability across a range of spatial and temporal frequencies. In this chapter, I compare the various relevant timescales for dynamic contrast stability and show how continuous WFS&C relaxes WFE requirements by at least an order of magnitude compared to static telescope control. Addressing the temporal aspect of WFE stability with the PASTIS model at its core, I draw connections to other studies going on in parallel with my PhD that treat dynamical WFE tolerancing.

Appendix A: Wavefront sensing and control on a segmented mirror telescope testbed

Appendix A describes my work on the JWST Optical Simulation Testbed (JOST), on which I was testing wide-field WFS&C. In particular, this chapter shows results from the fine-phasing of the segmented deformable mirror on several field points, performed with the LAPD algorithm.

Appendix B: System-level demonstrations of high contrast on the HiCAT segmented-aperture testbed

Appendix B presents the background and motivation of the HiCAT testbed and some of the work I contributed to reach a contrast of 10^{-8} . This chapter contains information about the overall testbed design, the different testbed modes it deploys, and a timeline overview over the various WFS&C techniques that were implemented during my time with the team.

2

A METHOD FOR ANALYTICAL TOLERANCING OF SEGMENTED TELESCOPES

The central question of my thesis is how to design, build and control segmented telescopes in combination with coronagraphy with the goal to image Earth-like extrasolar planets. The team at the Makidon laboratory first phrased this task in terms of a forward problem: make some assumptions about the telescope, add aberrations to its segments and model what kind of results that gives us in terms of coronagraphic performance and contrast. Reformulating this forward modeling process was the stepping stone for the analytical tolerancing model presented in this chapter, and it is only after we started inverting this new forward model that the applications for sensitivity studies became obvious.

The PASTIS forward model was originally developed by Lucie Leboulleux during her PhD (Leboulleux 2018), graduating in December 2018. She presented the problem inversion in the context of the eigenmodes of the optical system in Leboulleux et al. (2018b), which was the starting point of my own work on this topic. My first contribution to the subject was an improvement of the way we build the optical matrix that represents the contrast sensitivity of the segments. I presented this as a conference talk at SPIE Optics & Photonics in San Diego, California (USA) in August of 2019. The conference proceedings I wrote on that occasion (Laginja et al. 2019) served as a first draft for the peer-reviewed journal paper in this chapter.

The core question we are trying to answer in this paper is: If we set a mean contrast we want to observe in our dark hole, what are the wavefront error limits that we can tolerate on the telescope? This question is traditionally answered with WFE requirements expressed globally, over the entire telescope pupil. As a novel contribution, I introduce a statistical level into this problem that now allows us to analytically obtain independent segment WFE tolerancing limits. Based on an idea by Laurent Pueyo, the result turns out to confirm basic intuition in which certain segments have more influence on the coronagraphic contrast than others, and this paper quantifies it with a simple matrix inversion. Applying this analysis to the APLC coronagraphs of LUVOIR A, I show how the WFE requirements on certain segments in the pupil can be significantly relaxed, which directly informs the mechanical design process of large, segmented telescopes.

This chapter describes a general, theoretical framework to analytically calculate statistical WFE tolerances. We originally thought of segment-level tolerancing as the main result of the paper, but over time we identified more applications for coronagraph sensitivity analyses which have already been tackled during the course of my PhD. I directly contributed to these works that are based on the theory developed in this chapter, and I present some of them in Chap. 5.

Analytical tolerancing of segmented telescope co-phasing for exo-Earth high-contrast imaging

I. Laginja, R. Soummer, L. M. Mugnier, L. A. Pueyo, J.-F. Sauvage, L. Leboulleux, L. E. Coyle, J. S. Knight

Published in
JATIS, 7, 015004 (2021)

ABSTRACT

This paper introduces an analytical method to calculate segment-level wavefront error tolerances in order to enable the detection of faint extra-solar planets using segmented-aperture telescopes in space. This study provides a full treatment of the case of spatially uncorrelated segment phasing errors for segmented telescope coronagraphy, which has so far only been approached using ad hoc Monte-Carlo simulations. Instead of describing the wavefront tolerance globally for all segments, our method produces spatially dependent requirement maps. We relate the statistical mean contrast in the coronagraph dark hole to the standard deviation of the wavefront error of each individual segment on the primary mirror. This statistical framework for segment-level tolerancing extends the Pair-based Analytical model for Segmented Telescope Imaging from Space (PASTIS), which is based uniquely on a matrix multiplication for the optical propagation. We confirm our analytical results with Monte-Carlo simulations of end-to-end optical propagations through a coronagraph. Comparing our results for the Apodized Pupil Lyot Coronagraph designs for the Large UltraViolet Optical InfraRed (LUVOIR) telescope to previous studies, we show general agreement but we provide a relaxation of the requirements for a significant subset of segments in the pupil. These requirement maps are unique to any given telescope geometry and coronagraph design. The spatially uncorrelated segment tolerances we calculate are a key element of a complete error budget that will also need to include allocations for correlated segment contributions. We discuss how the PASTIS formalism can be extended to the spatially correlated case by deriving the statistical mean contrast and its variance for a non-diagonal aberration covariance matrix. The PASTIS tolerancing framework therefore brings a new capability that is necessary for the global tolerancing of future segmented space observatories.

2.1 Introduction

Imaging Earth-like exoplanets and searching for biomarkers is one of the key science objectives in space astronomy for the next decade. The close proximity of such planets to their host star, as well as a flux ratio on the order of 10^{-10} at visible wavelengths makes this a challenging task. These two goals can be achieved by using large-aperture telescopes for large light collecting areas and high angular resolution, in combination with static and dynamic starlight suppression techniques with coronagraphs and wavefront sensing and control (WFS&C) methods.

The invention of the coronagraph (Lyot 1939) synthesized with early ideas for the direct imaging of planets (Roman 1959) have led to several space mission concepts being developed toward this goal today. The Habitable Exoplanet Observatory (Gaudi et al. 2019) (HabEx) and the Large UV Optical InfraRed Surveyor (The LUVOIR Team 2019; Bolcar 2019) (LUVOIR) are two space-based concepts recently studied by NASA as possible future flagship missions. Their primary science objective is the direct detection and spectral characterization of habitable Earth-like planets (Roberge et al. 2019) and the search for life; they require primary mirror diameters of 4–15 meters. Meanwhile, the ground-based community is preparing for the era of extremely large telescopes (ELTs) where 30–40 meter class telescopes like the Thirty Meter Telescope (Simard et al. 2016) (TMT), the Giant Magellan Telescope (Fanson et al. 2018) (GMT) and the European Extremely Large Telescope (Ramsay et al. 2020) (E-ELT) will be equipped with coronagraphs and extreme adaptive optics systems to search for and characterize exoplanets (Guyon et al. 2012).

What unites all of these observatories is that they have significantly larger primary mirrors than their respective space-based and ground-based predecessors. This poses a number of problems that need to be solved, including considerations about overall mass, cost, and plausible launch vehicles for space-based missions. The logical consequence to this is that most of these observatories will have segmented primary mirrors, much like the Keck telescope (Mast et al. 1982) or the James Webb Space Telescope (JWST) (Acton et al. 2012; Perrin et al. 2018). This will allow for lighter-weight backplanes and foldable structures for launch purposes, or even in-situ space assembly (Polidan et al. 2018; Bowman et al. 2018).

Telescope segmentation introduces additional diffraction effects in the focal plane (Lightsey & Chrisp 2003; Yaitskova et al. 2003; Troy & Chanan 2003; Itoh et al. 2019), as well as sources for wavefront errors (WFE) due to segment misalignments and lighter mirror structure deformation in the form of localized segment-level aberration modes. All wavefront errors degrade the imaging performance in a high-contrast system (Crossfield & Troy 2007; Yaitskova & Troy 2011) as they generate light residuals all over the focal plane. Such WFE will directly impact the performance of the coronagraph instrument. There are a number of coronagraph designs that were developed specifically to maximize performance on telescopes with arbitrary apertures, which includes secondary obscurations, spiders and segmentation gaps (Sivaramakrishnan & Yaitskova 2005; Martinez et al. 2008; Soummer et al. 2009; Guyon et al. 2014; Zimmerman et al. 2016; Ruane et al. 2017).

All high-contrast instruments that aim at very high-contrast such as what is necessary to detect Earth-like planets will deploy strategies that combine static coronagraph masks in pupil and focal planes with active control of the electric field (Borde & Traub 2006; Pueyo et al. 2009; Mazoyer et al. 2018a,b) in order to create a zone of deep contrast in the final image plane, the dark hole (DH). To enable such wavefront control techniques, several methods for focal plane wavefront sensing have been developed (Groff et al. 2016; Jovanovic et al. 2018) to feed into a whole system of sensors and control loops that constitute the high-contrast instrument.

Even after careful cophasing of the segmented aperture and implementation of WFS&C techniques that reach the required star attenuation level, there will always be some residual errors due to drifts in the system (e.g., from thermal instabilities). These changes to the mechanical structure and in the optical train will have a direct effect on the observability of a faint point source, as a sufficient signal-to-noise ratio is needed for detection within confidence limits (Lyon & Clampin 2012). As a consequence, these high-contrast goals with segmented apertures impose severe requirements not only on static wavefront quality, but also stability requirements on the WFE as well as the overall mechanical structures of the telescope. There are various works that have tried to quantify these wavefront stability requirements for high-contrast imaging, both with and without segmented apertures in mind. The Nancy Grace Roman

Space Telescope (formerly known as the Wide-Field Infra-Red Survey Telescope, WFIRST) is a 2.4 meter monolithic space telescope with a large central obscuration and six thick, non-radial support struts (Krist et al. 2015) that render high-contrast imaging particularly challenging (Nemati et al. 2017a). Bound to launch in 2025, it will provide technology demonstrations for stellar coronagraphy at 10^{-9} contrast levels (Savransky & Garrett 2015) with the Roman Space Telescope Coronagraphic Instrument (CGI). Going to a segmented telescope introduces an increased number of degrees of freedom that will influence the final contrast. While there are solutions that aim to maintain a good contrast in the dark hole across integration times by means of continuous WFS&C (Pogorelyuk & Kasdin 2019), the problem must also be approached from an overall engineering perspective (Feinberg et al. 2017; Stahl 2017; East et al. 2018). In particular, the direct effects of segmentation on the final coronagraphic contrast (Yaitskova & Dohlen 2002; Stahl et al. 2015) are of interest in the context of high-contrast imaging, and there is an ongoing effort to characterize and quantify the requirements for such ultra-stable telescopes (Coyle et al. 2018, 2019b; Pueyo et al. 2019; Hallibert et al. 2019). Studies performing Monte-Carlo (MC) end-to-end (E2E) simulations (Moore & Redding 2018; Juanola-Parramon et al. 2019b) have confirmed the strict WFE requirements of a couple of tens of picometers over tens of minutes to enable the search for faint extra-solar planets, and analytical methods for the derivation of coronagraphic performance specifications have been proposed (Nemati et al. 2017b, 2020).

One thing that all of these studies have in common is that they define global WFE tolerances over the entire telescope pupil, where the segments have a random contribution to the overall aberrations. In this paper, we focus on analytically defining requirements on a segment-to-segment basis instead, using the Pair-based Analytical model for Segmented Telescope Imaging from Space (PASTIS), (Lebouilleux et al. 2017b, 2018a,b) which models the dark hole average contrast of a coronagraph on a segmented telescope as a function of the segment aberrations. We first introduce a new semi-analytical (SA) calculation method for the PASTIS matrix (Laginja et al. 2019). Then we show how to compute the statistical mean of the contrast using the PASTIS modes and extend the model inversion to calculate segment-level WFE requirements for a given target contrast.

A full error budget that aims at maintaining a particular DH contrast will contain WFE contributions both from spatially correlated as well as uncorrelated segments in the telescope pupil. The impact of aberrations made of correlated segments on coronagraph contrast has been studied in various cases, for example low-order Zernike modes as well as high-frequency checkerboard-like patterns in the pupil (N'Diaye et al. 2016; Moore & Redding 2018; Douglas et al. 2018). Aberrations made of spatially uncorrelated segments on the other hand have so far mostly been addressed in end-to-end simulations where the segments' amplitudes had equal standard deviations (Juanola-Parramon et al. 2019b; Stahl et al. 2015; Nemati et al. 2017b). In this paper we tackle the uncorrelated contribution, and establish analytically how to allocate WFE contributions to all segments individually. This addresses an essential component in the overall error budget, which had not been formally established yet. We then use the PASTIS approach to also generalize this to the correlated case.

In Sec. 2.2 we recall the development of the analytical propagation model and how the underlying PASTIS matrix was initially built through an analytical calculation. We then generalize the matrix calculation to all coronagraphs as well as segmented apertures with an extension to the semi-analytical matrix calculation, which eliminates the post-calibration step that used to be performed on a perfect coronagraph model. We show that the average contrast is always a quadratic function of the aberrations and drop the requirement of having a symmetrical dark hole. In Sec. 2.3 we perform the model inversion and validate the semi-analytical matrix. Further, we show that the PASTIS modes can be used to define a statistical framework for the analysis, additionally to their deterministic relation to the dark hole contrast. In Sec. 2.4 we derive the statistical mean contrast and its variance from two separate components that describe the imaging properties of the coronagraphic instrument on the one hand, and the thermo-mechanical segment statistical correlations on the other hand. We proceed with the calculation of independent segment-based WFE requirements and how to validate them in a statistical sense, and we show how to apply this to correlated segments. All simulations in Sec. 2.2-2.4 are done with a narrow-angle Apodized Pupil Lyot Coronagraph (APLC) on the primary pupil of LUVVOIR-A (see Fig. 2.1 and Sec. 2.5) at a wavelength of 500 nm, however, these methods can be applied to any combination of coronagraph and segmented tele-




Figure 2.1: *Left:* LUVUOIR-A design aperture with a diameter of 15 m. *Middle:* Narrow-angle apodizer for the LUVUOIR-A APLC, intended for exoplanet characterization. It uses a focal plane mask (FPM) with a radius of $3.5 \lambda/D$ (with λ the wavelength and D the telescope diameter). *Right:* Resulting coronagraphic image, with a dark hole from 3.4 to $12 \lambda/D$ and an average normalized intensity of 4.3×10^{-11} , which is the coronagraph floor in the absence of optical aberrations.

scope. A full demonstration of the PASTIS analysis is given in Sec. 2.5, where we calculate these segment tolerances for the case of three different APLC designs for LUVUOIR-A and highlight some consequences of this approach. In Sec. 2.6 we discuss our results and compare them to previously derived requirements and how they can be used in observatory error-budgets, and in Sec. 2.7 we report our conclusions.

Note that the main metric of the PASTIS model is the spatial average raw contrast in the dark hole (normalized coronagraphic intensity to peak of direct image), which is what we refer to as “contrast” throughout this paper, as opposed to a spatially dependent quantity. We also want to point out how we differentiate between this spatially averaged dark hole intensity, the “average DH contrast” on the one side, and a statistical mean (expectation value) of this averaged contrast over many optical propagations on the other side, the statistical “mean contrast”.

2.2 PASTIS model of telescope segment-level aberrations in high-contrast coronagraphy

The PASTIS model was initially established for a perfect coronagraph using an analytical propagation model for aberrated pairs of segments (Lebouilleux et al. 2018b); the application to real coronagraphs required a second-step numerical calibration. Here, we generalize the model to any coronagraph on any segmented aperture geometry by using a semi-analytical derivation of the PASTIS matrix. We also show that the validity of the PASTIS results is not limited to symmetrical dark holes, but extends to non-symmetrical ones as well. Independently of the way the PASTIS matrix is calculated (analytically or semi-analytically), the derivations and conclusions that we build on the PASTIS approach retain their analytical power and potential.

2.2.1 Matrix formalism to calculate the average dark hole contrast

The goal of PASTIS is to model coronagraphic images in the presence of optical aberrations on a segmented primary mirror, which can be represented for example by using localized Zernike polynomials. This basis is an obvious possible choice since segment-level piston, tip/tilt, focus and astigmatism are naturally occurring aberrations from segment misalignments, for example in three-mirror anastigmat (TMA) designs such as JWST (Acton et al. 2004, 2012) or LUVUOIR (The LUVUOIR Team 2019; Juanola-Parramon et al. 2019b). Although beyond the scope of this paper, the PASTIS approach can also be applied directly to any other function basis, for example to represent mirror wavefront errors induced by thermo-mechanical effects (mounting, backplane deformations, etc.) (Coyle et al. 2019a; East et al. 2019; Wells & East 2019). In this paper, we simply expand the phase aberration in the segmented pupil ϕ_s as a sum of local (segment-level)




Figure 2.2: Piston pair aberrations on a segmented pupil (top) and the resulting image plane intensity distributions in the dark hole (bottom), using the narrow-angle APLC for LUVOIR-A in an E2E propagation model. The left three panels show different interference pairs with corresponding Young-like interference fringes, while the right panel shows a random distribution of local piston on all segments of the pupil and the resulting image plane intensity. All plots appear on the same scale.

Zernike polynomials (Leboulloux et al. 2018b, Eq. 9):

$$\phi_s(\mathbf{r}) = \sum_{(k,l)=(1,0)}^{(n_{seg},n_{zer})} a_{k,l} Z_l(\mathbf{r} - \mathbf{r}_k), \quad (2.1)$$

where \mathbf{r} is the pupil plane coordinate, ϕ_s the phase from the segmented primary and n_{seg} is the total number of segments, indexed by k . The $a_{k,l}$ are the Zernike coefficients with Noll index (Noll 1976) l up to the maximum Zernike n_{zer} , and Z_l is the l^{th} Zernike. In this paper, we limit the study to a single Zernike mode (piston; index $l = 0$) as illustrated in Fig. 2.2. Hence we drop the l index in all consecutive equations, but the PASTIS methodology is applicable to any Zernike mode, combination thereof, or other types of segment-level modes.

In high-contrast coronagraphy the best contrast is not typically obtained for the perfect aperture without any aberration, but more commonly in the presence of a wavefront control solution using deformable mirrors (Pueyo et al. 2019; Mazoyer et al. 2018a). Therefore, we are studying the response of the coronagraphic system to a perturbation around that solution. Defining ϕ_{DH} as the phase solution for best DH contrast and ϕ_s as the segmented perturbation, the phase can be divided into

$$\phi = \phi_{DH} + \phi_s. \quad (2.2)$$

High-contrast coronagraphy requires exquisite wavefront quality around the dark hole solution, and therefore we assume the small aberration regime for ϕ_s , where the electric field $E(\mathbf{r})$ is well approximated as an affine function of the phase: $E(\mathbf{r}) = P(\mathbf{r}) e^{i\phi(\mathbf{r})} \simeq P'(\mathbf{r}) + i\phi_s(\mathbf{r})$. The phase $\phi_s(\mathbf{r})$ is zero where the pupil aperture $P(\mathbf{r})$ is zero, and $P'(\mathbf{r})$ is a complex pupil that includes the wavefront solution to produce the static dark hole (with both phase and amplitude contributions, and including static errors). Note that the phase ϕ_{DH} is not necessarily small (Mazoyer et al. 2018a,b).

Using Fourier optics for a scalar description of the electric field and of its propagation, the coronagraph propagation can be represented by a linear operator C . This is a valid assumption for Lyot-style coronagraphs, for example an APLC (Soummer et al. 2003a; N'Diaye et al. 2015b, 2016) such as the one illustrated in Fig. 2.1 for the LUVOIR-A coronagraph design, or a vortex coronagraph (Foo et al. 2005; Mawet et al. 2013). High-order vortex designs would need a special treatment for the specific low-order modes (e.g.,

defocus) they reject perfectly (Ruane et al. 2017), but the tolerancing of such global modes is not the main purpose of PASTIS anyway. We can hence express the intensity distribution in the final image plane as

$$I(\mathbf{s}, \phi) = |C\{P'\}(\mathbf{s}) + iC\{\phi_s\}(\mathbf{s})|^2, \quad (2.3)$$

with \mathbf{s} the image plane coordinate. This intensity is therefore the sum of three terms (Leboulleux et al. 2018b, Eq. 16):

$$I(\mathbf{s}) = |C\{P'\}|^2 + 2\Re\{C\{P'\}C\{\phi_s\}^*\} + |C\{\phi_s\}|^2. \quad (2.4)$$

In most cases of interest, we will be working in a symmetrical DH. It can be shown that the spatial average of the linear cross-term in Eq. 2.4 over a symmetrical DH is zero (Leboulleux et al. 2018b, Appendix A). This simplifies Eq. 2.4 to a quadratic function of the phase:

$$\langle I(\mathbf{s}) \rangle_{DH} = \langle |C\{P'\}|^2 \rangle_{DH} + \langle |C\{\phi_s\}|^2 \rangle_{DH}. \quad (2.5)$$

The main metric used in this paper is the spatial average contrast over the extent of the dark hole, $\langle \dots \rangle_{DH}$, so by using $c'_0 = \langle |C\{P'\}|^2 \rangle_{DH}$, we can express the average dark hole intensity as:

$$\langle I(\mathbf{s}) \rangle_{DH} = c'_0 + \langle |C\{\phi_s\}|^2 \rangle_{DH}. \quad (2.6)$$

Using the expression for the phase decomposition from Eq. 2.1 in Eq. 2.6, we can derive the intensity as a function of all aberrated segment pair combinations:

$$\langle I(\mathbf{s}) \rangle_{DH} = c'_0 + \langle |C\left\{\sum_k^{n_{seg}} a_k Z(\mathbf{r} - \mathbf{r}_k)\right\}|^2 \rangle_{DH}, \quad (2.7)$$

and therefore:

$$\langle I(\mathbf{s}) \rangle_{DH} = c'_0 + \sum_i^{n_{seg}} \sum_j^{n_{seg}} a_i a_j \langle C\{Z(\mathbf{r} - \mathbf{r}_i)\} C\{Z(\mathbf{r} - \mathbf{r}_j)\}^* \rangle_{DH}. \quad (2.8)$$

This double sum combines all pairs of segments where segment i has an aberration amplitude a_i of the localized phase aberration $Z(\mathbf{r} - \mathbf{r}_i)$. These cross-terms from each aberrated pair of segments are very similar to Young interference fringes, and this forms the basic idea behind the PASTIS model (Leboulleux et al. 2018b). The orientation and periodicity of these fringes depend on the separation and orientation of the according aberrated pair, as displayed in Fig. 2.2.

It is important to note that the pair-wise model is not an ad-hoc idea to build the model by pairs. It derives from the fact that we expand the primary mirror phase on a discrete number of segments. Since we build a propagation model for the intensity, the “pairs” simply appear in Eq. 2.8 from all the cross-terms when calculating the square modulus of the electric field in Eq. 2.7.

Eq. 2.8 can be readily re-written as a matrix multiplication:

$$c = c'_0 + \mathbf{a}^T M \mathbf{a}, \quad (2.9)$$

where c is the average contrast in the dark hole, c'_0 the coronagraph floor (i.e., the average contrast in the dark hole at best contrast with ϕ_{DH} , in the absence of phase perturbations), M is the PASTIS matrix with elements m_{ij} , \mathbf{a} is the aberration vector of the local Zernike coefficients on all discrete n_{seg} segments and \mathbf{a}^T its transpose. The elements of the PASTIS matrix M in Eq. 2.9 therefore directly identify as:

$$m_{ij} = \langle C\{Z(\mathbf{r} - \mathbf{r}_i)\} C\{Z(\mathbf{r} - \mathbf{r}_j)\}^* \rangle_{DH}. \quad (2.10)$$

While this derivation is always true in the most common case of a symmetrical DH, there are coronagraph designs that produce half-sided dark holes (Por 2020b). We can show that the quadratic dependency

of the contrast on the phase perturbations remains true in this most general case. We rewrite Eq. 2.4 in a similar matrix form as Eq. 2.9, but preserving the linear term:

$$c = c'_0 + \mathbf{v}^T \mathbf{a} + \mathbf{a}^T M \mathbf{a}, \quad (2.11)$$

where \mathbf{v} is a vector that does not need to be expressed explicitly here. If we take the derivative of this equation and solve for the aberration vector \mathbf{a}_0 that provides the minimum contrast c_0 , we can identify $\mathbf{a}_0 = -M^{-1}\mathbf{v}/2$ and $c_0 = c'_0 - 1/4\mathbf{v}^T M^{-1}\mathbf{v}$, and therefore eliminate the linear term by performing a simple change of variable:

$$c = c_0 + (\mathbf{a} - \mathbf{a}_0)^T M (\mathbf{a} - \mathbf{a}_0). \quad (2.12)$$

This quadratic expression is similar to Eq. 2.9, but with a segmented mirror perturbation solution $\mathbf{a}_0 \neq 0$ that improves contrast compared to the case without aberrations. As discussed above, we also assume a wavefront control solution with deformable mirrors to be included in the term $P'(\mathbf{r})$ (and hence c_0), with both amplitude and phase contributions. Therefore this guarantees that the best contrast in the presence of that wavefront control solution and DH is obtained for $\mathbf{a}_0 = 0$, which in turn means that any arbitrary segment aberration vector \mathbf{a} will always degrade the contrast. Note that this does not preclude to have a non-zero static segmented correction included as part of the term ϕ_{DH} . This is equally true in broadband light: When summing over wavelengths, the quadratic nature of Eq. 2.12 remains true, albeit with different coefficients c_0 , a_0 and M .

We have shown that the average dark hole contrast is always a quadratic function of a segmented phase perturbation ϕ_s , which can be discretized into a per-segment aberration amplitude vector \mathbf{a} , coefficients on a modal basis. We can calculate this average dark hole contrast for any aberration vector directly, using the PASTIS matrix expression (Eq. 2.9). This is particularly interesting and efficient since it does not require end-to-end optical simulations, and only involves simple linear algebra. Furthermore, this analytical expression can be inverted to establish a segment-level wavefront error budget that meets a given level of contrast. This will be detailed in Sec. 2.4.

2.2.2 Semi-analytical calculation of the PASTIS matrix

The PASTIS matrix M can be calculated using the original analytical approach for a perfect coronagraph, then calibrated numerically for a real coronagraph and to include pupil features (e.g., support structures) (Lebouilleux et al. 2018b, Eq. 20). This approach was validated against an end-to-end model for the 36-segment ATLAST telescope pupil with an APLC (Lebouilleux et al. 2018b, Fig. 7) to within an error of 3%.

Here, we introduce another way to calculate the PASTIS matrix using an end-to-end simulation (Laginja 2020) of the average dark hole contrast for all individually aberrated segment pairs, from which we can identify semi-analytically the matrix elements in Eq. 2.10. This presents the advantage of enabling a direct calculation of the matrix for any telescope geometry, any coronagraph, and any choice of segment-level aberrations (including fully numerical ones such as segment figures induced by thermo-mechanical effects).

The phase for each segment pair is expressed as a Zernike aberration:

$$\phi_{ij}(\mathbf{r}) = a_i Z(\mathbf{r} - \mathbf{r}_i) + a_j Z(\mathbf{r} - \mathbf{r}_j). \quad (2.13)$$

We denote by $c_{ij} = \langle I_{ij}(\mathbf{s}) \rangle_{DH}$ the average dark hole contrast, on the pair of segments i, j , $\phi_{ij}(\mathbf{r})$, that can be calculated numerically for a small wavefront aberration and compared to the quadratic expression of Eq. 2.6 under the linear expansion of this phase term:

$$\begin{aligned} c_{ij} &= c_0 + \langle |a_i C\{Z(\mathbf{r} - \mathbf{r}_i)\} + a_j C\{Z(\mathbf{r} - \mathbf{r}_j)\}|^2 \rangle_{DH} \\ &= c_0 + a_i^2 \langle |C\{Z(\mathbf{r} - \mathbf{r}_i)\}|^2 \rangle_{DH} + a_j^2 \langle |C\{Z(\mathbf{r} - \mathbf{r}_j)\}|^2 \rangle_{DH} \\ &\quad + a_i a_j 2 \langle C\{Z(\mathbf{r} - \mathbf{r}_i)\} C\{Z(\mathbf{r} - \mathbf{r}_j)\}^* \rangle_{DH}. \end{aligned} \quad (2.14)$$

The elements m_{ij} of the PASTIS matrix M (Eq. 2.10) can then be identified directly in Eq. 2.14 as:

$$c_{ij} = c_0 + a_i^2 m_{ii} + a_j^2 m_{jj} + 2a_i a_j m_{ij}, \quad (2.15)$$

where the diagonal terms of the PASTIS matrix are

$$m_{ii} = \frac{c_{ii} - c_0}{a_i^2}, \quad (2.16)$$

and the off-diagonal elements:

$$m_{ij} = \frac{c_{ij} + c_0 - c_{ii} - c_{jj}}{2a_i a_j}. \quad (2.17)$$

For simplicity, we choose the same calibration aberration amplitude $a_c = a_i = a_j$ for both segments. Throughout our analytical development above, a_i is in units of radians, as ϕ is a phase. Since the PASTIS matrix can be normalized to any units though, the units of the aberration amplitude a_c can be chosen freely in the computation of Eqs. 2.16 and 2.17. The units of the PASTIS matrix are therefore in “contrast per square of units of a_c ” (contrast having no physical dimension), which is consistent with Eq. 2.9. Note that in the presented case in Fig. 2.3, the units of the aberration amplitude a_c is waves. The aberration amplitude a_c has to be chosen such that the global pupil aberration it results in yields an average DH contrast higher than the contrast floor, but small enough to remain in the small phase aberrations linear regime. This will be discussed further in Sec. 2.2.3. The matrix is symmetric by definition since $c_{ij} = c_{ji}$. Off-diagonal elements m_{ij} of the PASTIS matrix (Eq. 2.17) can be negative, which is not an issue since the only constraint for the matrix is to be positive semi-definite to ensure positive eigenvalues, since they correspond to each mode’s contrast. This will be discussed in detail in Sec. 2.3.

We could potentially calculate the matrix elements m_{ij} (Eq. 2.10) directly by calculating those complex electric field quantities. Usually though, full end-to-end simulators that calculate the image plane intensity are readily available and necessary for multiple other reasons. This means that choosing to calculate the PASTIS matrix through image plane intensities makes it more flexible and portable to other simulators. More importantly, working with intensities allows us to measure an empirical PASTIS matrix without the estimation errors and computational overheads of using an electric field estimator, allowing this theory to be experimentally tested.

In summary, the PASTIS matrix is constructed in two steps: (1) Calculate aberrated images I_{ij} and their corresponding dark hole average contrast c_{ij} for each pair of aberrated segments i, j , and (2) use these contrast values to identify analytically the elements of the PASTIS matrix M based on Eqs. 2.16 and 2.17. Here, the numerical calculation of these aberrated images for pairs of segments using an end-to-end simulator (see Chap. 3.1) replaces the analytical expression of Young fringes between pairs of segments (Lebouilleux et al. 2018b). This approach provides more accuracy, flexibility and generality for use with any coronagraph and telescope geometry, since the analytical approach has to be calibrated using a numerical simulation anyway.

2.2.3 Validating the semi-analytical PASTIS matrix

The semi-analytical PASTIS matrix for the narrow-angle LUVOIR APLC is calculated following Sec. 2.2.2 and shown in Fig. 2.3.

The PASTIS matrix shows how some segments have a higher impact on the final contrast than others. This is visible along the diagonal, which records the contrast contribution from each individual segment alone. For example, segments 65-120 have a lower contrast contribution, as they correspond to the darker areas of the apodizer on the outer two rings of the aperture (see Fig. 2.3, right panel). This effect is also visible on the innermost ring of hexagons. We can also notice streaks of negative values in the matrix in the off-axis areas, as discussed in Sec. 2.2.2.

We validate the semi-analytical PASTIS matrix by comparing the PASTIS contrast obtained with the matrix formalism of Eq. 2.9 to the contrast from an E2E simulator using the same inputs. We show the comparison in Fig. 2.4. The coronagraph floor for this particular APLC design in the absence of aberrations is 4.3×10^{-11} . The PASTIS model starts to diverge from the E2E calculation at large WFE root-mean-square (RMS) where the linear approximation of the phase breaks down. Note that the choice of $a_c = 1/500$ of the wavelength (used in the presented example, and is on the order of 1 nm in visible) on a single segment yields a global pupil WFE RMS of 1.67×10^{-5} waves, which translates into an average DH contrast just above the coronagraph floor, but keeps it in the small aberration regime. Here, the accuracy of the semi-analytical




Figure 2.3: Semi-analytical PASTIS matrix of the 120 segment LUVOIR-A design with the narrow-angle APLC (*left*). The matrix is symmetric by construction and the dark streaks correspond to negative values. The diagonal elements show which segments have more impact on the contrast than others. The outermost segments (65–120) have lower matrix values because of the darker apodization for these segments. This is clearly visible in the superimposed image of the apodizer on the segments (*right*).

matrix approach is significantly higher than that of the fully analytical matrix because the construction of the PASTIS matrix is based on the actual E2E simulation as opposed to a post-calibrated analytical fringe model.

2.3 Model inversion and statistical mean contrast derivation

Once the PASTIS matrix has been calculated, Eq. 2.9 gives a fully analytical expression of the dark hole average contrast for any random segment-level aberration \mathbf{a} . This makes PASTIS particularly well suited for error budgeting analyses, compared to otherwise computation-intensive Monte-Carlo analyses. More interestingly, this analytical model can be inverted to determine the pupil plane aberration vector \mathbf{a} that meets a specific average contrast target c_t , using an eigendecomposition of the PASTIS matrix. We also show that the model inversion to obtain the target contrast as a function of eigenmodes is achieved both in a deterministic and statistical sense.

2.3.1 Eigendecomposition of the PASTIS model and mode-segment relationship

The PASTIS matrix M is square and symmetric by construction, and therefore diagonalizable. We perform the eigendecomposition:

$$M = UDU^T, \quad (2.18)$$

where U is unitary, hence invertible and $U^{-1} = U^T$. The columns of U are the eigenmodes of the PASTIS matrix M , which can be written as column vectors \mathbf{u}_p of $U = (\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_p, \dots, \mathbf{u}_{n_{modes}})$ and n_{modes} is the total number of eigenmodes (which is equal to the total number of segments n_{seg}), indexed by p . D is a diagonal matrix whose diagonal elements are the eigenvalues λ_p of the matrix M ; it is the diagonalized PASTIS matrix D . The analysis of the eigenmodes \mathbf{u}_p provides information about the critical modes of the system that can be used to place tolerances on segment cophasing and stability. The full set of modes of the LUVOIR-A primary with the narrow-angle APLC is shown in Fig. 2.5, and a selection of modes in Figs. 2.7, 2.8, and 2.9. The eigenvalues λ_p shown in Fig. 2.6 indicate how much each mode contributes to the final image contrast if applied to the pupil in their natural normalization, without any imposed weighting. This figure shows that the high-spatial frequency modes (to the left) have a much higher impact than the lower-spatial frequency modes (to the right).

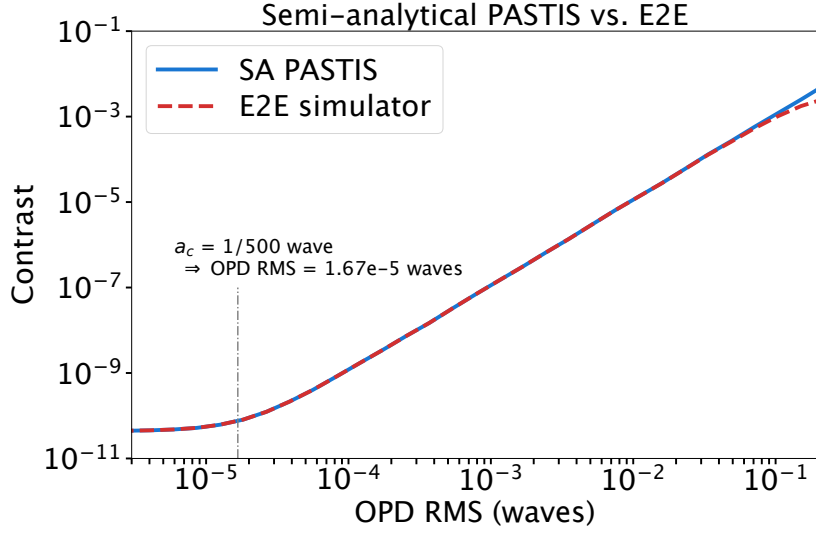


Figure 2.4: Average dark hole contrast as a function of wavefront error using both an end-to-end simulator (dashed red) and the PASTIS matrix propagation (solid blue). In a hockey stick graph behavior, the contrast is limited by the coronagraph itself at low wavefront errors corresponding to the flattened out curve to the left (at c_0). From about 10^{-4} waves to 10^{-1} waves of WFE RMS, the contrast is limited by segment phasing aberrations. In this range the estimation error of PASTIS is 0.06% compared to the reference E2E model. A calibration aberration per segment a_c of 1/500 wave, on a 120 segment pupil, translates to a global WFE of 1.67×10^{-5} waves when calibrating the PASTIS matrix diagonal, or 2.36×10^{-5} with two simultaneously aberrated segments, which is in the small aberration regime of the model, just above the coronagraph floor. The curves shown are obtained as the mean of the same 20 random realizations for each RMS value, both for the E2E simulator and the PASTIS propagation. At large wavefront errors (close to 0.1 waves RMS) the linear approximation breaks down and the two curves no longer match perfectly.

The PASTIS modes \mathbf{u}_p form an orthonormal basis set that allows us to express any arbitrary, segment-based pupil plane aberration \mathbf{a} as a linear combination of the modes \mathbf{u}_p with mode weighting factors b_p :

$$\mathbf{a} = \sum_{p=1}^{n_{modes}} \mathbf{u}_p b_p. \quad (2.19)$$

This can also be written as:

$$\mathbf{a} = \mathbf{U} \cdot \mathbf{b}, \quad (2.20)$$

indicating a basis transformation between the mode basis and the segment basis. The inverse basis transformation is thus given by $\mathbf{b} = \mathbf{U}^{-1}\mathbf{a}$. This relationship demonstrates the physical equivalence of working in the mode-basis or in the segment basis, as we can transform any expression in one space into an expression of equivalent meaning in the other space. We further explore the physical meaning of the PASTIS modes in Sec. 2.3.2 and 2.3.3.

2.3.2 Contrast as a function of the eigenmodes

The mode weights \mathbf{b} will depend on how much each individual mode contributes to the final contrast, and their associated eigenvalue. Inserting Eq. 2.20 into Eq. 2.9 allows us to define this relationship:

$$\begin{aligned} c - c_0 &= (\mathbf{U}\mathbf{b})^T \mathbf{M} (\mathbf{U}\mathbf{b}) \\ &= \mathbf{b}^T \mathbf{U}^T \mathbf{M} \mathbf{U} \mathbf{b} \\ &= \mathbf{b}^T \mathbf{D} \mathbf{b}, \end{aligned} \quad (2.21)$$

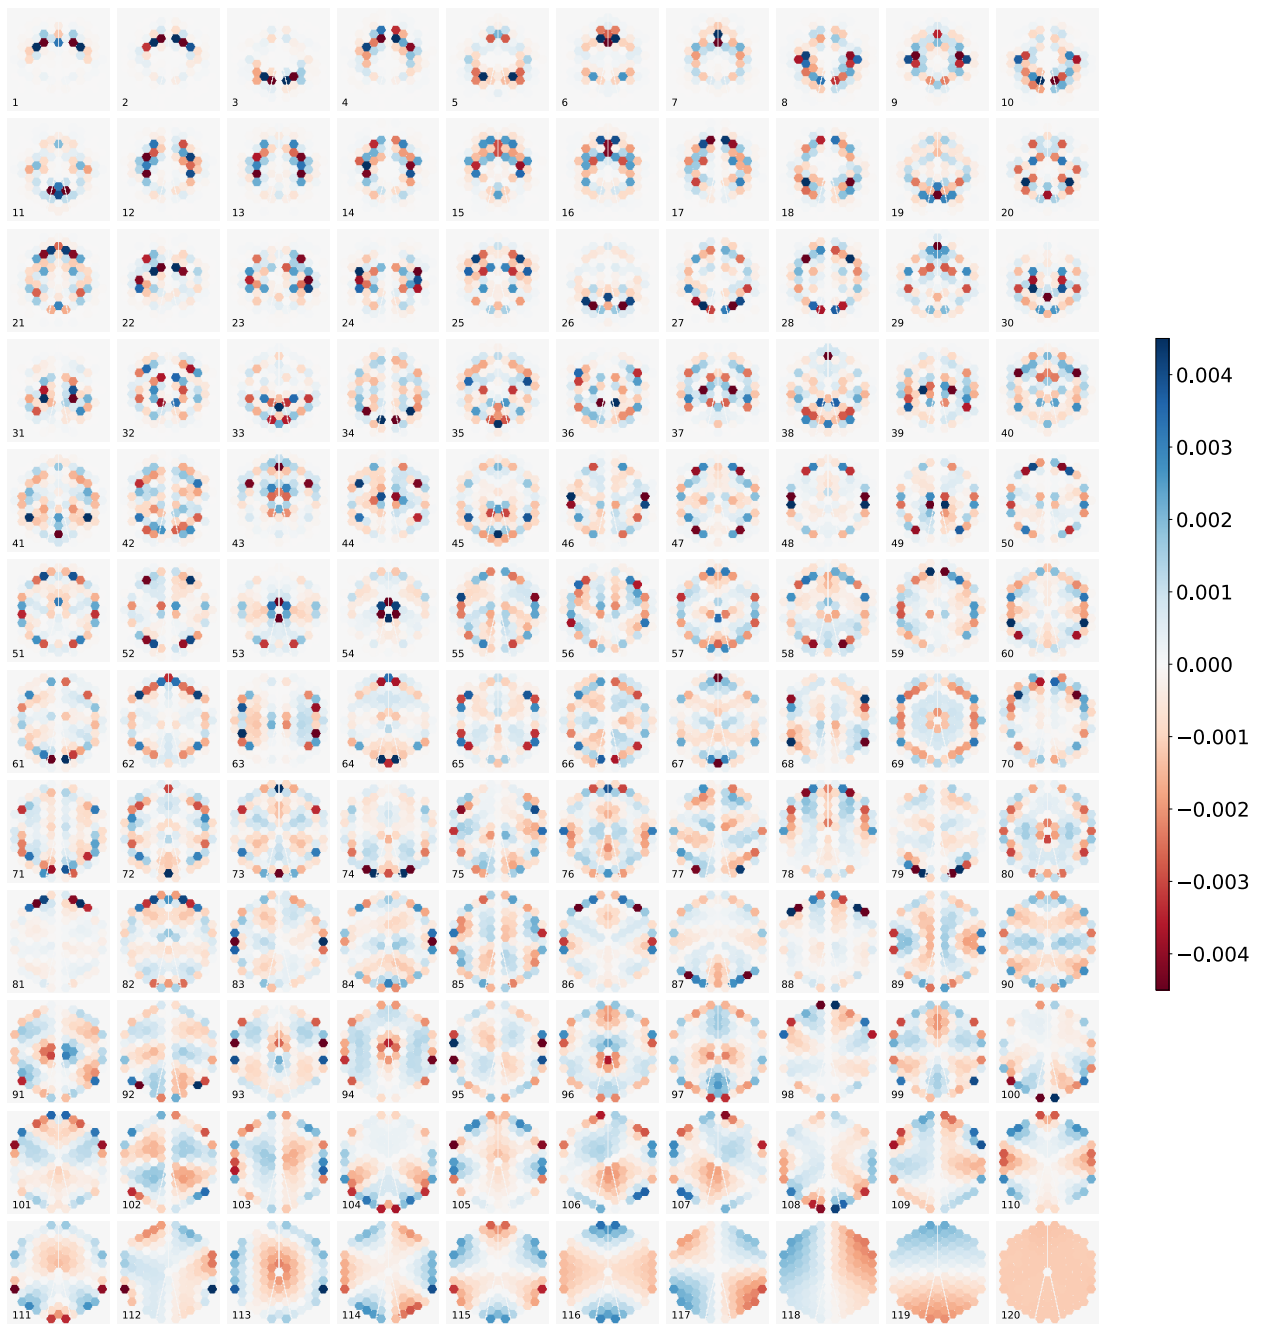


Figure 2.5: All PASTIS modes for the LUVOIR-A narrow-angle APLC, for local piston aberrations, sorted from highest to lowest eigenvalue. The modes are unitless, showcasing the relative scaling of the segments to each other, and between all modes. They gain physical meaning when multiplied by a mode aberration amplitude b_p in units of wavefront error or phase. Their respective eigenvalues and hence relative impact on final contrast is displayed in Fig. 2.6.




Figure 2.6: Eigenvalues, or sensitivity of contrast to mode index p , for the piston PASTIS modes of the LUVOIR-A telescope with the small FPM coronagraph design, shown in Fig. 2.5. Note how the PASTIS matrix and modes do not depend on the target contrast, but they do on the choice of telescope geometry and coronagraph, making them the proper modes of the optical system.

and finally:

$$c - c_0 = \sum_p^{n_{modes}} b_p^2 \lambda_p. \quad (2.22)$$

The final contrast is therefore the sum of all squared mode weights, multiplied by their respective eigenvalue. Since the modes contribute independently to the final contrast (they are orthonormal by construction), we can define a per-mode contrast as:

$$c_p = b_p^2 \lambda_p, \quad (2.23)$$

and obtain that the total contrast is the sum of all individual contrast contributions:

$$c = c_0 + \sum_{p=1}^{n_{modes}} c_p. \quad (2.24)$$

We can then find the p -th mode weight that gives the allocated contrast contribution c_p as:

$$b_p = \sqrt{\frac{c_p}{\lambda_p}}. \quad (2.25)$$

Eq. 2.25 gives the weighting factor for each PASTIS mode when it has a particular contrast contribution c_p . We can illustrate this expression by calculating the mode weights corresponding specifically to a uniform contrast contribution of the overall target contrast c_t over all modes, $c_p = (c_t - c_0)/n_{modes}$. Then we calculate them as

$$\tilde{b}_p = \sqrt{\frac{c_t - c_0}{n_{modes} \cdot \lambda_p}}, \quad (2.26)$$

where the \tilde{b}_p is the particular set of mode weights in the case of a uniform contrast allocation across all modes. The resulting mode weights $\tilde{\mathbf{b}}$ for a total contrast allocation of $c_t = 10^{-10}$ are shown in Fig. 2.10.

In Fig. 2.11 we confirm the validity of the mode weights \tilde{b}_p by showing the average dark hole contrast from an end-to-end propagation of the cumulative wavefront error for all modes. The linearity of the plot, as well as the end value at the target contrast c_t validates the uniform contrast allocation to each PASTIS mode from Eq. 2.26.




Figure 2.7: Low-impact modes with high tolerances for the narrow-angle APLC on the LUVOIR-A telescope, for local piston aberrations. These modes have little impact on the final contrast - they are similar, but not equal, to discretized Zernike modes and the coronagraph rejects them very well by design.




Figure 2.8: Mid-impact modes with medium tolerances for the narrow-angle APLC on the LUVOIR-A telescope, for local piston aberrations. These modes have medium impact on the final contrast, relatively speaking. These modes show mostly low spatial frequency features except for high spatial frequency components in the parts of the pupil where the apodizer covers most of the segments.




Figure 2.9: High-impact modes with low tolerances for the narrow-angle APLC on the LUVOIR-A telescope, for local piston aberrations. These modes have the highest impact on the final contrast. They consist entirely of high spatial frequency components in the parts of the pupil where the apodizer (and other pupil plane optics) are the most transmissive.




Figure 2.10: PASTIS mode weights for the uniform contrast allocation across all modes. The low-index modes to the left, which correspond to high spatial frequencies, have a lower WFE tolerance than the low spatial frequency modes with high index to the right. These mode amplitudes are inversely proportional to the eigenvalues associated with each mode (Eq. 2.26), and they scale the modes such that each of them contributes the same contrast c_p to the overall target contrast. The cumulative contrast response of the modes multiplied by these weights is shown in Fig. 2.11.




Figure 2.11: Cumulative contrast from all PASTIS modes when allocating the total contrast uniformly across all modes (see Sec. 2.3.2). For instance, the measured contrast corresponding to the first 60 accumulated weighted modes is about 0.7×10^{-10} . We multiply all modes by their respective mode amplitude \tilde{b}_p and propagate them cumulatively to the image plane, both with the E2E simulator (dashed red) and the PASTIS propagation (solid blue). Without any of the modes applied, we get the contrast floor from the coronagraph c_0 , while application of all modes together yields the requested target contrast, here $c_t = 10^{-10}$. Each mode is allocated an equal contrast contribution c_p to the final contrast, which results in a linear cumulative contrast curve. Note how neither line starts at the coronagraph floor because the lowest-index mode already adds a contrast contribution on top of the baseline contrast. The corresponding PASTIS mode weights to obtain this uniform allocation of contrast per mode, is shown in Fig. 2.10.

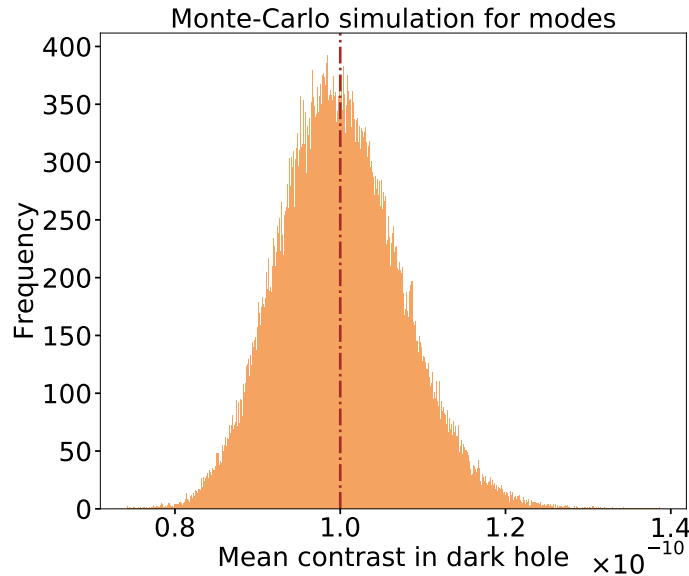


Figure 2.12: Validation of the uniform contrast allocation across PASTIS modes with an E2E Monte-Carlo simulation, drawing random sets of mode weights from a normal distribution with standard deviations σ (Eq. 2.28), corresponding to an equal contrast allocation per mode (Eq. 2.26) with weights $\tilde{\mathbf{b}}$. The overall WFE of one realization is the sum of all weighted modes for each set, and is propagated in an E2E simulation. The histogram represents 100,000 realizations of the average dark hole contrast for a target contrast of 10^{-10} .

2.3.3 Statistical mean of the contrast from mode amplitudes

In this section, we analyze the properties of the model in the statistical sense to prepare a framework for the segment-level error budget in Sec. 2.4. We extend the formalism from purely deterministic mode weights \mathbf{b} to random variables. We obtain the statistical mean contrast by substituting Eq. 2.23 into Eq. 2.24 and taking the mean:

$$\langle c \rangle - c_0 = \sum_p^{n_{modes}} \langle b_p^2 \rangle \lambda_p. \quad (2.27)$$

Assuming zero-mean normal distributions of the PASTIS modes, we can readily identify their variance as

$$\sigma_p^2 = \langle b_p^2 \rangle. \quad (2.28)$$

We verify this for the uniform contrast allocation per mode (Eq. 2.26) in an E2E Monte-Carlo simulation where we draw random samples of the PASTIS mode coefficients $\tilde{\mathbf{b}}$, following zero-mean normal distributions \mathcal{N} with standard deviations σ_p , according to the uniform contrast allocation from Eq. 2.26:

$$\mathbf{b} = (\mathcal{N}(0, \sigma_1), \mathcal{N}(0, \sigma_2), \dots, \mathcal{N}(0, \sigma_p)). \quad (2.29)$$

Fig. 2.12 shows the result of such a Monte-Carlo simulation where we sum each randomly weighted, individual set of modes to a unique wavefront map, propagate it to the image plane with the E2E simulator and measure the spatial average contrast in the dark hole. We validate that the statistical mean of all these contrast values is the target contrast c_t for which we calculated the vector of mode standard deviations σ in the first place, in this case 10^{-10} .

In summary, PASTIS provides an analytical model to go from a set of segment aberrations to the DH average contrast in the coronagraphic image. The calculation of the PASTIS matrix eigenmodes allows to invert this model: we can set a target contrast and allocate WFE amplitudes to each of the system's eigenmodes to reach that target contrast. Since they form an orthonormal set of modes, their individual contrast contributions add independently, each with its own sensitivity, as a fraction of the average contrast in Sec. 2.3.2, on top of their statistical description in Sec. 2.3.3. These optical modes contain the full

information of the image formation system, including the apodizer, coronagraph, primary geometry and other optical components in the system, and are distinct from and unaware of the mechanical behavior of the telescope. By this nature, we can use them to understand the fundamental limitations for high contrast imaging with a segmented aperture, which will in sequence be further constrained by thermo-mechanical properties of the telescope as we describe in Sec. 2.4.

2.4 Segment-level tolerance statistics

Sec. 2.3 details how we can construct mode-based requirements that satisfy a target on the spatial average contrast in the image plane. This relationship is not only valid in a deterministic sense (relating a mode amplitude directly to its contrast contribution), but also in the statistical sense (relating the standard deviation of random modes to the overall dark hole contrast). The PASTIS modes form an orthonormal basis where each mode contributes independently to the dark hole average contrast: the contrast of a sum of weighted modes is equal to the sum of the contrast contributions for all weighted modes. The PASTIS matrix M and its eigenmodes fully describe the optical propagation through the system in terms of WFE effects from a segmented aperture on the average contrast in the DH with a coronagraph. However, they do not contain any information about thermal or mechanical effects necessary to describe the final performance of a given segmented observatory. A sound framework to develop error budgets on segmented apertures therefore requires the combination of both the optical response of the telescope and coronagraph (encapsulated in the PASTIS modes), and the thermo-mechanical response of the telescope and the observatory (encapsulated in the segment aberration covariance matrix).

The goal of this section is to combine the information about the imaging formation through the coronagraph on a segmented mirror with the thermo-mechanical properties of the observatory, in order to establish requirements to reach a target DH contrast.

2.4.1 Statistical mean contrast and its variance in segmented coronagraphy

We can calculate the statistical mean contrast of the DH spatial average directly from Eq. 2.9, exploiting the fact that the trace of a scalar is the scalar itself, and that $\text{tr}(AB) = \text{tr}(BA)$:

$$\begin{aligned} \langle c \rangle &= c_0 + \langle \mathbf{a}^T M \mathbf{a} \rangle = c_0 + \langle \text{tr}(\mathbf{a}^T M \mathbf{a}) \rangle \\ &= c_0 + \langle \text{tr}(M \mathbf{a} \mathbf{a}^T) \rangle = c_0 + \text{tr}(M \langle \mathbf{a} \mathbf{a}^T \rangle), \end{aligned} \quad (2.30)$$

and finally:

$$\langle c \rangle = c_0 + \text{tr}(M C_a), \quad (2.31)$$

where C_a is the $n_{seg} \times n_{seg}$ segment covariance matrix, containing the as-built thermo-mechanical correlations between segments. Eq. 2.31 allows us to calculate the statistical mean of the average DH contrast directly from the knowledge of the segment covariance matrix, no matter if there is correlation between the segments or not, combining the imaging properties of the high contrast imaging system, contained in M , with the thermo-mechanical behavior of the instrument contained in C_a .

Similarly, we can derive an analytical expression for the variance $\text{Var}(c)$ of the DH contrast. Assuming that a follows a zero-mean Gaussian distribution, the variance for Eq. 2.9 takes the very simple form (Rencher & Schaalje 2008, Theorem 5.2c):

$$\text{Var}(c) = 2 \text{tr}[(M C_a)^2]. \quad (2.32)$$

These two equations provide an unambiguous closed form derivation of the mean contrast and its variance from the optical model of the imaging system (encapsulated in the matrix M), and from the thermo-mechanical properties of the telescope (captured by the matrix C_a). The PASTIS matrix M knows nothing of the thermo-mechanical effects of the observatory and is obtained by diffractive modeling of the coronagraph, while the segment covariance matrix comes from thermal and mechanical modeling of the observatory and is completely detached from the image formation system of the telescope. The two matrices together (M and C_a) fully describe the statistical response of the coronagraph system to a

particular WFE allocation on segments and therefore allow to establish a set of top-level requirements on segment tolerances for an observatory.

The enabling aspect of Eqs. 2.31 and 2.32 for segment-level tolerancing is that the trace is invariant under a basis transformation. It follows that if either one of the two matrices, the PASTIS matrix or the thermo-mechanical covariance matrix, is expressed in its diagonal basis, the expressions for the contrast mean and variance simplify greatly, as we show in the following sections. The segment tolerancing can thus be achieved either by diagonalizing M , or by doing so with C_a . We have treated the case of diagonalizing the PASTIS matrix M in Sec. 2.3, where we describe the analytical framework for segmented telescope tolerancing in the diagonal basis that most naturally describes the optical sensitivity of the system to the DH contrast. In the following two sections, we turn to a basis that diagonalizes the segment covariance matrix instead, permitting us to perform the tolerancing on appropriate system modes.

2.4.2 Uncorrelated segment-level requirements

In finding a diagonal basis for the thermo-mechanical matrix, the easiest case is when C_a is already diagonal, which physically corresponds to independent segments on the primary mirror. In this case, the diagonal elements of C_a , namely the segment variances $\langle a_k^2 \rangle$, fully describe the effect of the primary mirror segments on the DH contrast, and the statistical mean of the contrast in Eq. 2.31, $\langle c \rangle$, finds a simple expression similar to Eq. 2.27:

$$\langle c \rangle = c_0 + \sum_k^{n_{seg}} m_{kk} \langle a_k^2 \rangle. \quad (2.33)$$

Similarly to the mode-based error budget presented in Sec. 2.3, we now want to find a segment-based error budget to formulate the WFE limits on each segment that reach a specific statistical mean target contrast $c_t = \langle c \rangle$. Turning to a statistical mean contrast allows us to define a similar allocation of contrast contributions to all segments as we did statistically (and deterministically) in the PASTIS mode basis (Eq. 2.24). The most straightforward way of doing this is to allocate the target contrast equally to all segments:

$$\langle a_k^2 \rangle m_{kk} = \frac{\langle c \rangle - c_0}{n_{seg}}. \quad (2.34)$$

If we define μ_k as the standard deviation of the WFE on the k -th segment, comparably to Eq. 2.28:

$$\mu_k^2 = \langle a_k^2 \rangle, \quad (2.35)$$

then by combining the three previous equations we obtain the per-segment WFE requirement for this particular contrast allocation:

$$\mu_k^2 = \frac{\langle c \rangle - c_0}{n_{seg} m_{kk}}. \quad (2.36)$$

The expression in Eq. 2.36 lets us calculate a per-segment requirement for all individual segments in the pupil of a coronagraphic instrument, given a statistical mean target contrast. The main assumption for this is that the segments are independent from each other, and that we have access to the statistical mean value of the contrast. While the mean contrast is easily measurable on images through averaging, this might not be the case for an ultra-stable facility like LUVOIR. However, the statistical mean contrast is an important quantity to perform segment-level WFE tolerancing, especially with regards to mirror manufacturing.

In a more physical sense, we know that the intensity or contrast is proportional to the variance of the WFE. Therefore, the total final contrast over the full pupil is proportional to the sum of the segment variances, which is also proportional to the sum of the contrast for all segments, conforming with Parseval's theorem. We validate this independent segment-level error budget for three different LUVOIR coronagraphs in Sec. 2.5.

2.4.3 Case of correlated segments

While the assumption of statistically independent segments brings insights into WFE tolerances of segmented mirrors for coronagraphic imaging, it is not general enough to encompass all possible modes for such telescopes where large-scale thermo-mechanical drifts occur (for example, backplane “flapping” mode around the folding motion of the primary mirror). Here we discuss extensions of the PASTIS approach to the case of correlated segments.

In the case of independent segments, the covariance matrix C_a of the aberration vector \mathbf{a} is a simple diagonal matrix holding the segment variances μ_k^2 . However, the covariance matrix C_a is no longer diagonal for correlated segments, because of mechanical coupling for example due to large-scale backplane deformations. In this case, the statistical mean contrast and its variance remain analytically computable with Eqs. 2.31 and 2.32. The tolerancing can be done by performing an eigendecomposition on C_a , which will diagonalize it and provide an orthonormal set of eigenmodes that describe the mechanical perturbations of the telescope system, which is also known as the Karhunen-Loève basis. By writing $C_a = VC_{th}V^T$, we obtain the PASTIS matrix in this new basis, M' , through the transformation matrix V , where $M = VM'V^T$ ($m'_{kk} \in M'$) still describes the optical properties of the system. In this basis, Eq. 2.31 takes the form :

$$\langle c \rangle = c_0 + \text{tr}(M'C_{th}). \quad (2.37)$$

Since the thermo-mechanical covariance matrix C_{th} is diagonal, we identify its diagonal elements as the thermo-mechanical mode variances s_k^2 . Like for Eqs. 2.27 and 2.33, this simplifies Eq. 2.37 yet again to:

$$\langle c \rangle = c_0 + \sum_k^{n_{th}} m'_{kk} s_k^2, \quad (2.38)$$

which allows us to make a reasonable allocation of contrast contributions across all n_{th} thermo-mechanical eigenmodes. In the same way like Eq. 2.36 calculates a per-segment variance in the basis of independent segments, we can use Eq. 2.38 to tolerance the per-mode variances s_k^2 to any given target contrast, albeit this time for individual mechanical eigenmodes. Similarly, a transformed expression can be found for the contrast variance in Eq. 2.32:

$$\text{Var}(c) = 2 \text{tr}[(M'C_{th})^2]. \quad (2.39)$$

In this most general case of correlated segments, the knowledge of the segment-level covariance matrix C_a and its diagonal eigenbasis C_{th} supersedes the simpler description in terms of segment-level variances that is only possible in the uncorrelated case (Sec. 2.4.2). It allows us to express the mean contrast explicitly as a function of variances of thermo-mechanical modes that can be tolerated in a similar fashion to what was done in Sec. 2.3.

We have presented a quantitative, fully analytical method to calculate segment-level tolerances for a high-contrast instrument on a segmented aperture telescope. These follow directly from the PASTIS matrix for which we provided a new, semi-analytical way for its calculation that exploits a numerical simulator to compute the effects of the segments on the intensity in the image plane. We encode the optical and thermo-mechanical properties of the observatory separately, with the PASTIS matrix M and the segment covariance matrix C_a , which when put together allow for the analytical calculation of the expected mean contrast and its variance. These equations are invariant under a basis transformation, which permits us to find an appropriate diagonalized basis in order to derive individual WFE tolerances. This can either be done by diagonalizing M , as we showed in Sec. 2.3, or by finding a diagonal basis for C_a . A special case is given if C_a is naturally diagonal due to independent segments on the segmented mirror; in this case, we derive a per-segment requirement map by following the analytical framework set forth in Sec. 2.3. In the more general case of correlations between the segments, due to thermo-mechanical properties of the telescope, we diagonalize C_a and use the same analysis principles in the Karhunen-Loève basis of C_a , which allows us to calculate per-mode WFE tolerances. While the deformation matrix C_a will be acquired through thermo-mechanical modeling and can include thermal, vibrational or gravitational perturbations, the PASTIS matrix M and the PASTIS modes give insight into the purely optical properties of the observatory, and the sensitivity of the optical system to contrast.

In the next section, we validate the segment-level error budget in the case of independent segments (Sec. 2.4.2) for three different APLC designs for LUVUOIR.

2.5 Application to LUVOIR WFE tolerancing

The LUVOIR study (The LUVOIR Team 2019) has two point-design cases (LUVOIR-A and LUVOIR-B), respectively 15 m and 8 m in diameters, each designed as a TMA and containing a suite of scientific instruments that include coronagraphs. The LUVOIR-A coronagraphic instrument (Pueyo et al. 2017) includes a suite of three numerically optimized APLC coronagraphs (Por et al. 2020) with focal plane mask diameters that maximize the exo-Earth yield in both detection and characterization (Stark et al. 2015, 2019). The optical train of an APLC (Lebouilleux et al. 2018b, Fig. 5) contains an apodizer in the pupil plane that modulates the optical beam in amplitude, a focal plane mask occulting the on-axis point-spread-function (PSF) core, and a Lyot stop in the subsequent pupil plane that blocks the light diffracted at the focal plane mask (FPM) (Soummer et al. 2003a; N'Diaye et al. 2015b, 2016). Of the three LUVOIR APLCs, we used the smallest FPM coronagraph, or narrow-angle coronagraph, for theory validation in the previous sections. It is typically used for spectroscopic characterization in the wavelength band where molecular oxygen and water can be detected ($0.76 \mu\text{m}$ and $0.94 \mu\text{m}$). Planet detection can however be performed at shorter wavelengths (e.g., around $0.5 \mu\text{m}$) where a given angular size corresponds to a larger inner working angle in diffraction resolution units (λ/D). This larger inner working angle corresponds to a larger FPM, and a larger FPM allows for apodizer designs with a higher throughput and a more robust coronagraph design, which is where the trade-off between the three designs (narrow-, medium- and wide-angle) is made. The three LUVOIR-A APLC designs are shown in the top row of Fig. 2.13. The corresponding FPM have radii of 3.50 , 6.82 and $13.38 \lambda/D$ respectively, followed by a hard edge annular Lyot stop, whose inner and outer diameters are 12.0% and 98.2% of the circumscribed diameter of the apodizers. The resulting coronagraphic dark hole sizes are $3.4\text{--}12$, $6.7\text{--}23.7$ and $13.3\text{--}46.9 \lambda/D$ respectively, with a coronagraph floor c_0 of 4.3×10^{-11} , 3.9×10^{-11} and 3.9×10^{-11} for the three designs.

In this section, we present a full analysis to obtain segment requirements for these three LUVOIR APLCs, and validate the results by performing Monte-Carlo simulations with an E2E simulator. We also take a deeper look into the narrow-angle APLC by analyzing the PASTIS mode-based decomposition of the individual-segment requirements. This monochromatic analysis was performed at a wavelength of 500 nm , which is the lower limit wavelength for the LUVOIR coronagraphs and where we expect to detect planets.

2.5.1 Segment requirements and Monte-Carlo simulations for three APLC designs

We first calculate the PASTIS matrix for each of these three APLC designs, according to the methodology described in Sec. 2.2.2. We can then establish a segment-level error budget in the assumption of uncorrelated segments, according to Eq. 2.36. In Fig. 2.13, we show the resulting segment requirement maps for a target contrast of $c_t = 10^{-10}$ and for all three APLC designs.

It is important to note that these requirement maps do not represent the WFE over the segmented pupil, but instead show the standard deviations on the tolerable WFE for each segment in order to retrieve, statistically, the desired mean target contrast. In this sense, the maps in Fig. 2.13 are a prescription for the drawing of random segment WFE realizations like the examples shown in Fig. 2.14. These random maps are then propagated with the E2E simulator and their average contrast values build the MC histograms in Fig. 2.15. One big takeaway point from Fig. 2.13 is that the segment requirements are not uniform across the pupil, but clearly follow the apodization of the coronagraph mask. The PASTIS matrix holds knowledge of the optical effect of not only the segments but also the coronagraph instrument on the final contrast, so by including that knowledge into the derivation of the segment constraints we obtain a requirement map optimized for that particular instrument. Moreover, we can observe a direct trade-off between the coronagraph apodization and the per-segment requirements - the more aggressive the apodization and the lower the throughput, the more we can relax the requirements on the more concealed segments within one coronagraph. However, more aggressive apodization usually comes with smaller FPM coronagraphs that filter low-order modes less, which will lead to more stringent overall requirements. This leads to a direct trade-off between FPM size, throughput and WFE requirements (see also Sec. 2.6).

These requirement maps can be calculated for any target contrast in the range of validity of the PASTIS model, which we discussed in Sec. 2.2.3. We can verify Eq. 2.36 by running MC simulations with the E2E




Figure 2.13: *Top:* The three baseline apodizer designs for LUVOIR-A, a narrow-angle (left), medium-angle (middle) and wide-angle (right) mask (details see Sec. 2.5). *Bottom:* Segment tolerance maps for narrow-angle (left), medium-angle (middle) and wide-angle (right) APLC designs on LUVOIR-A for a target contrast of $c_t = 10^{-10}$, at a wavelength of 500 nm. All three tolerance maps are shown on the same scale. Note how each segment value denotes the standard deviation of a zero-mean normal distribution from which the segment aberrations in WFE RMS are drawn. The minimum and maximum values of these maps are, from left to right: 7 and 116 pm, 25 and 93 pm, and 92 and 181 pm.




Figure 2.14: Four random segment-based WFE maps drawn from a zero-mean normal distribution and the per-segment standard deviations from the left prescription map in Fig. 2.13, for the narrow-angle APLC design and a target contrast of 10^{-10} . After each random map is created, we propagate it through the end-to-end simulator and record its average contrast to build the left MC simulation in Fig. 2.15.

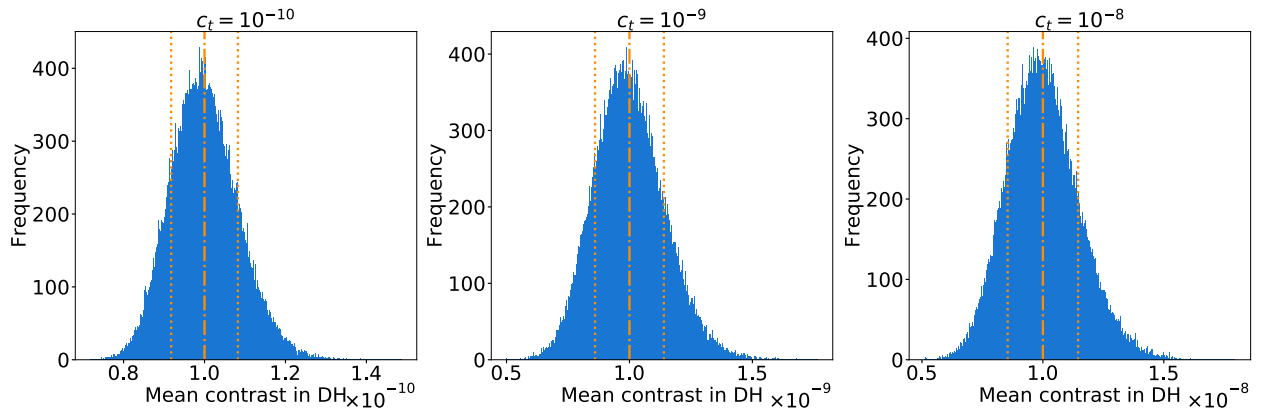


Figure 2.15: Validation of the independent segment tolerancing with E2E Monte-Carlo simulations, for different target contrasts, using the narrow-angle APLC design. Each segment k in one of the 100,000 WFE realizations is drawn from a zero-mean normal distribution with standard deviation μ_k . The dashed-dotted lines mark the target contrast of each case, which are successfully recovered by the mean values of the histograms, in accordance with their analytical calculation in Eq. 2.31. The dotted lines mark the 1-sigma confidence limits of this contrast distribution, which are 8.3×10^{-12} , 1.4×10^{-10} and 1.5×10^{-9} for the three target contrasts 10^{-10} , 10^{-9} and 10^{-8} respectively, and they accord with the numbers calculated by Eq. 2.32.

simulator across a grid of different coronagraph instruments and target contrasts. Using a range of target contrasts $c_t = 10^{-10}$, 10^{-9} , and 10^{-8} on the narrow-angle baseline LUVOIR APLC design, we first calculate the segment constraints (the requirement maps for 10^{-9} and 10^{-8} are not shown here, but they show the same spatial distribution over the segments as in Fig. 2.13, only different by a proportionality factor). We draw the WFE amplitude for each individual segment from a zero-mean normal distribution and its standard deviation μ_k :

$$\mathbf{a} = (\mathcal{N}(0, \mu_1), \mathcal{N}(0, \mu_2), \dots, \mathcal{N}(0, \mu_k)). \quad (2.40)$$

We use these random aberration amplitudes on all segments to compose a WFE map on the segmented pupil and then propagate this WFE map through the E2E simulator to measure the resulting spatial average contrast in the dark hole. Doing this 100,000 times for each target contrast case, we obtain the histograms shown in Fig. 2.15. The mean of the resulting MC simulations clearly recovers the target contrast for which the segment requirements have been calculated, which is indicated by the dashed-dotted line. Both these mean values, as well as the standard deviations, indicated with the dotted lines in Fig. 2.15, agree with the theoretical values calculated analytically from the segment covariance matrix (Eq. 2.31 and Eq. 2.32). Also, we have verified the correct recovery of the same range of target contrasts by means of MC simulations for the other two APLC designs shown in Fig. 2.13 (resulting histograms not shown in this paper).

2.5.2 Modal analysis of the segment-based requirements

The segment requirement maps were obtained assuming a uniform contrast allocation across all segments (Eq. 2.34). We also assumed statistically independent segments, so that their correlation matrix C_a was diagonal. Here, we further explore this uniform error budget in the segment basis by analyzing the corresponding distribution of proper system modes of the optical system, the PASTIS modes. Using the transformation matrix U from the eigendecomposition of the PASTIS matrix M , we can calculate the corresponding covariance matrix in the PASTIS mode basis with $C_b = U^T C_a U$. Given this linear transformation, if the covariance matrix is diagonal in one space, we do not expect it to be diagonal in the other space. The matrix C_b , obtained from the diagonal segment covariance matrix C_a assembled from the requirement map, is illustrated on the left hand side in Fig. 2.16. This figure also compares the extracted standard deviations for PASTIS modes along the diagonal of C_b , with the PASTIS mode weights previously calculated in Sec. 2.3, on the right hand side. Although C_b is not diagonal, this is a legitimate comparison. The PASTIS matrix M is always diagonal in its own eigenbasis, expressed as matrix D in Sec. 2.3.1. This is why the average contrast expression from the statistical mode weights in Eq. 2.27 only requires the diagonal elements of




Figure 2.16: *Left:* Covariance matrix C_b , calculated from the diagonal covariance matrix in segment-space, C_a , with $C_b = U^T C_a U$. Although there are clearly some correlations present between the high-index PASTIS modes in the top right corner (low spatial frequencies), this does not matter as long as we are in the PASTIS segment basis, where the PASTIS matrix is diagonal. When this is the case, the mean contrast only depends on the diagonal elements of C_b (Eq. 2.27). *Right:* PASTIS mode amplitudes for the case of independent segments in WFE RMS (solid red). They are extracted from the mode covariance matrix C_b , after constructing an error budget assuming independent segments that contribute equally to the total contrast, at 500 nm. Overlapping (dashed grey), we can see the mode weights from the uniform contrast allocation to all PASTIS modes from Fig. 2.10. We can clearly see how compared to that flat allocation, the independent-segment error budget increases the tolerances of low-index modes (left) that have less segment correlation, and dampens the tolerances of high-index modes (right) that are highly correlated, low-spatial frequency modes.

the covariance matrix C_b , no matter whether it is diagonal or not, i.e. whether the mode weights show some correlations or not. The difference with respect to the mode weights \tilde{b}_p (Eq. 2.26) obtained under the assumption of a uniform contrast allocation per mode (Fig. 2.10) is notable: the mode weights of low mode index have increased tolerances, which is very interesting from a system design point of view, while large index modes (above ~ 90) are strongly attenuated in the PASTIS mode basis error budget for independent segments. This is also clearly visible in Fig. 2.17 where the contrast contribution per mode is relatively flat at a low mode index, but drops to negligible contributions at high-index modes. For the case of the flat contrast allocation across modes, this figure shows a flat line at $(c_t - c_0)/n_{seg}$ for comparison (dashed line). This effect can be well understood by looking back at Fig. 2.5, where high-index modes appear to be very similar to low-order Zernike modes, therefore having highly correlated segments, and the low-index modes appear as high-spatial frequencies, i.e. with more uncorrelated segments. Therefore, the construction of a segment-level error budget for uncorrelated segments creates a modal distribution with extremely low weights on the PASTIS modes that have highly correlated segments (high-index modes), as seen in Fig. 2.16. Also, since the mode contrast contribution is directly related to the mode weight (Eq. 2.25), the same effect is visible in the allocated contrast per mode (Fig. 2.17).

To illustrate this further, we calculate a cumulative contrast plot similar to Fig. 2.11, which was initially obtained for a uniform contrast allocation per mode. The new result is shown in Fig. 2.18, where we can see that it is no longer linear, i.e. the modes no longer contribute equally to the total mean contrast. The slope of the blue curve is indicative of the allocated tolerances for each mode: the low-spatial frequency PASTIS modes on the right hand side now contribute significantly less to the final contrast, while the first ~ 80 modes contribute more, while still resulting in the exact cumulative target contrast. This is consistent with the behavior discussed in Figs. 2.16 and 2.17.

2.6 Discussion

The results we obtain for the three LUVUOIR APLC designs in Sec. 2.5, under the assumption of statistically independent segments, span more than one order of magnitude from the most constrained segment on




Figure 2.17: Contrast per individual PASTIS mode when derived from the error budget in which all segments contribute independently and equally to the final contrast (solid blue). High-index modes to the right, which correspond to low-spatial frequencies and therefore highly correlated segments (see Fig. 2.5), are highly attenuated and contribute negligible amounts to the contrast. The uniform contrast allocation across all modes at $(c_t - c_0)/n_{seg}$ is indicated with the dashed grey line (the contrast floor has been removed in both curves).




Figure 2.18: Cumulative contrast of the PASTIS modes for the error budget in which all segments contribute independently and equally to the final contrast (solid blue), compared to the case in Fig. 2.11 where all modes contribute the same contrast (dashed grey). The high-index modes have negligible contrast impact (see also Fig. 2.17) as they correspond to low-spatial frequency, highly correlated segments. This plot also confirms the assumption that the mode covariance matrix C_b is nearly diagonal.

the small APLC to the most relaxed segment on the large APLC (7 pm to 181 pm). These results show not only a dependency on the coronagraph design, but also a wide range of segment requirements within the pupil for one single coronagraph. The segment with the most stringent requirement can tolerate a standard deviation of 7 pm local piston error on the narrow-angle APLC for a target contrast of 10^{-10} , at a wavelength of 500 nm, which is comparable to previous results for segment-based piston errors on LUVUOIR-A (10 pm at a wavelength of 575 nm (Juanola-Parramon et al. 2019b)), while studies on other apertures and with other segment numbers quote similar numbers (Stahl et al. 2015; Nemati et al. 2017b). However, we show that segments in other parts of the pupil have a much higher local piston tolerance, a standard deviation up to 116 pm on the small APLC design, which suggests that not all segments need to be held to the same tolerance level. Instead, we can relax the segment-level requirements on those segments that do not influence the average DH contrast as much, while still obtaining the same statistical mean contrast. Local relaxation of the wavefront error limits on certain parts of the pupil can be exploited for example for the backplane mechanical design and observatory-level control strategy.

In particular, the tolerances will also depend on the total number of segments in the pupil. While this has not been studied systematically in this paper, PASTIS can enable such work. For example, with fewer segments in the aperture, the spatial frequencies corresponding to segment misalignments will be lower. Therefore, WFE from these misalignments will be more filtered by the coronagraph, which might lead to increased tolerances. Inversely, with more segments in the pupil, the highest spatial frequencies from segment misalignments will not be filtered by the FPM and might thus result in lower tolerances, as is already the case for the most sensitive PASTIS modes. Wavefront aberrations from a mirror with fewer segments will be typically more filtered by the coronagraph. However, more and smaller segments will introduce high spatial frequencies that will diffract light into the image beyond the outer working angle, decreasing the impact of each segment misalignment on the DH. Overall, the number of segments in the pupil will raise competing effects that will influence the overall tolerances, and will be highly influenced by the type of coronagraph (Laginja et al. 2020).

In Sec. 2.5.1, we briefly mentioned the observed trade-off between coronagraph throughput, FPM size and per-segment requirements. For a given coronagraph design, we observe relaxed requirements for those segments in the pupil that are more concealed by the apodizer (i.e., more black in the apodizer image). The overall tolerance over the entire pupil also increases with the size of the FPM. The larger the FPM, the higher the rejection in particular of low-order spatial modes, which correspond to high mode indices where the modes are similar to low-order, Zernike-like global modes. This results in higher mode weights, which becomes obvious in Fig. 2.10, where these high-index modes on the right side show higher WFE tolerances. When moving to the independent-segment error budget in Fig. 2.16, this effect becomes less obvious as the low-order mode tolerances get dampened due to our assumption of uncorrelated segments contributing equally to the mean contrast. The projection of these uncorrelated segments onto the mode basis favors high-spatial frequency modes. This leaves the low-order modes statistically weak, as they would otherwise contribute to inter-segment correlation. However, this is only true for these fully uncorrelated segment-level WFE contributions. This uncorrelated error budget is over-constraining the low-order (high-index) modes (see Fig. 2.16), therefore not taking advantage of the coronagraphic rejection of these spatially correlated modes. A complete error budget will need to allocate contrast contributions separately between the correlated and the uncorrelated components of the WFE. The final result will therefore have a modal weight distribution in-between the solid (fully uncorrelated) and dashed lines (uniform contrast across modes) in Fig. 2.16.

Additionally, large-FPM APLCs have higher throughput apodizers. This means that their larger coronagraphic rejection (associated with the larger FPM) contributes more to the WFE tolerance relaxation than the apodizer throughput itself. Therefore, a true optimization of the WFE tolerances will be a trade-off between the FPM size and the fraction of apodization in the pupil. Further, this tolerancing work introduces new design considerations for high-contrast instruments, which is the optimization of the coronagraphic component with respect to segment phasing tolerances. Such an optimization will aim to release the segment tolerances while keeping a reasonable contrast goal, with the ultimate goal to maximize exoplanet yield, which should be explored in future work.

In a realistic telescope of course, the segments are typically correlated due to the deformations of

the backplane structure. The development of technologies that support increased wavefront stability of segmented telescopes is actively being worked on today (Coyle et al. 2019b). This includes precise methods for thermo-mechanical modeling, and measuring of such deformation effects on the segmented primary mirror. The underlying segment correlations can be provided either as a segment covariance matrix, or in the form of thermo-mechanical system eigenmodes, in which case we can directly use Eq. 2.37 for a tolerancing analysis, after expressing the PASTIS matrix in this new basis.

While local piston errors have been shown to have the largest impact on the contrast (Nemati et al. 2017b; Juanola-Parramon et al. 2019b), we can generate a PASTIS matrix for other local Zernike modes as well (e.g., tip/tilt, focus, astigmatism). The feasibility of this has already been shown in the analytical approach (Lebouilleux et al. 2018b) and should hence be regarded as a mere functional addition. Moreover, if we have knowledge of telescope-design and hardware dependent local aberration modes (e.g., from effects such as adhesive shrinkage, bulk temperature, coatings, or gravity), these can be used as well to build the PASTIS matrix and derive their corresponding segment-level requirements. Instead of evaluating the tolerance levels mode by mode, we can also calculate a multi-mode PASTIS matrix that incorporates combinations of local modes (e.g., piston–tip–tilt, or combinations of custom modes) in order to derive segment tolerances that will take into account that more than one distinct local aberration mode is contributing to the overall WFE.

Another application of extended PASTIS matrices is the generalization to high-spatial frequency effects (e.g., from polishing). Instead of building a PASTIS matrix with pair-wise Zernike aberrations of segments, we can use sine waves locally on the segments in lieu of Zernikes. Each spatial frequency and orientation would then be a new local mode, and we can use many of them to make a multi-mode PASTIS matrix in the same way as with any other local modes. While a generalization to a continuous distribution of frequencies to build tolerances in terms of a Power Spectral Density (PSD) might be possible, it is beyond the scope of the present discussion¹. Nevertheless, the generalization to a few sine-wave frequencies (e.g. with their corresponding speckles localized at the inner and outer working angle, or in the middle of the DH) is a direct and straightforward extension of the present illustration and would provide meaningful input for tolerancing purposes of polishing errors.

The presented tolerancing model provides WFE limits on the segments, but it does not define how these limits are to be maintained. Relying purely on the mechanical stability of the telescope will not be enough to stay within these requirements and therefore an active optics system will be needed to measure WFE deviations and compensate for the residuals. Such an active optics system will include WFS&C as well as signal-to-noise considerations, as the wavefront sensor needs enough photons to provide an accurate wavefront estimate and correction (Pueyo et al. 2019).

Finally, the application of PASTIS to ground-based observatories is possible, but will have to take additional effects in account. While future large segmented telescopes (TMT, E-ELT, GMT) will reach contrast levels of 10^{-7} to 10^{-8} within the next decade (Kasper et al. 2008), which is sufficient to enter the high-contrast regime that the PASTIS model can be applied to, future work will have to include the effects of residual turbulence in order to truthfully represent coronagraphic observations on those observatories. As a first common approximation for this purpose, the coronagraphic PSF can be expressed as the sum of a static and a dynamic contribution (Ygouf et al. 2013), so the PASTIS analysis can be used for a characterization of the static part.

2.7 Conclusions

The goal of PASTIS, as established by Lebouilleux et al. (Lebouilleux et al. 2018b), is an analytical propagation model to calculate the average dark hole contrast in a coronagraphic system, in the presence of segment-level aberrations. This is achieved with a closed-form expression in Eq. 2.9 that depends exclusively on the PASTIS matrix M , acting on the aberration amplitudes on all segments. In this paper, we extended the calculation of the matrix M to a semi-analytical approach, where the optical propagation of segment aberrations is performed numerically before assembling the M matrix analytically (Eqs. 2.16

¹We show some preliminary results of PASTIS with more general aberrations of a continuous DM, obtained after the publication of this paper, in Appendix 2.A.

and 2.17). This makes the model more accurate as it includes all details of the optical system as provided with the end-to-end simulator. We also show that the model holds even for a non-symmetrical DH. The semi-analytical PASTIS approach is therefore a flexible tolerancing tool that can be adapted readily to any telescope geometry or coronagraph, as shown in Sec. 2.5. It can be used to study trade-offs between coronagraph designs that will provide certain tolerance distributions over the segments on the primary mirror, and telescope-level engineering constraints implemented in other parts of the observatory.

We used the model to derive analytical expressions for the statistical mean and variance of the average DH contrast (Eqs. 2.31 and 2.32). This opens the possibility for WFE tolerancing of a segmented observatory. In addition to the optical properties modeled by the PASTIS matrix M , these expressions involve the segment-level covariance matrix C_a that describes the thermo-mechanical properties of the telescope. Indeed, deformations of the backplane structure typically lead to correlated segment poses (e.g., backplane flapping modes around the folding motion of the primary mirror). It is this combination of the optical with the thermo-mechanical characteristics that lays the foundation to a complete and analytical method for the tolerancing of segmented aberrations.

The key to calculate WFE requirements with this framework is to find the diagonal basis of either of the two matrices, M or C_a . The PASTIS matrix M can easily be diagonalized by means of an eigendecomposition, in which case the tolerancing can be performed on the PASTIS eigenmodes. They form an orthonormal set of modes, representing the proper optical system modes of the given observatory and coronagraphic instrument, and allow to analyze its fundamental limitations in terms of WFE propagation in segmented aperture coronagraphy. These eigenmodes contribute additively to the image plane average contrast according to their mode-level tolerances, which correspond to the standard deviations per mode σ_p , associated with the statistical mean of the average DH contrast.

Additionally, if we have information on the thermo-mechanical behavior of the instrument or of the whole observatory embodied in the segment aberration covariance matrix (e.g., though finite-element simulations), we can choose to work in the eigenbasis of these thermo-mechanical perturbations, a.k.a. the Karhunen-Loève basis, to perform the tolerancing.

This allows us to put requirements on structural deformations that impact the segmented primary mirror (e.g., backplane and mirror support structures) and can be modeled as rigid-body motions at the segment level. A better approximation of these thermo-mechanical modes will be obtained with a multi-mode PASTIS matrix by combining multiple local Zernikes and/or ad-hoc local aberration modes, as described in Sec. 2.6.

In the simplified case of independent segments, the covariance matrix C_a contains only diagonal elements, representing the individual segment WFE variances μ_k^2 . We have built a segment-based error budget by allocating equal contribution to contrast from all segments. This allowed us to calculate WFE requirements for all segments individually, building segment requirement maps as we showed in Sec. 2.5 for the three APLC designs of the LUVUOIR-A telescope. The advantage of this method is two-fold: Firstly, rather than calculating WFE tolerances globally over the entire pupil, we can obtain a WFE standard deviation per segment, which can locally lead to a relaxation in requirements. Secondly, we do not need to perform full Monte-Carlo simulations that evaluate different realizations of wavefront error maps for that purpose, instead, we can calculate these requirements analytically in one single step.

The analysis presented in this paper is statistical but static, i.e. without temporal evolution; the extension to dynamical drift rates depends on the observing scenario and wavefront control strategy, which will put the PASTIS propagation model on different time scales (Coyle et al. 2019b; Pueyo et al. 2019). Future work will address such dynamic analysis methods for continuous wavefront sensing and control cases on ultra-stable telescopes.

Appendix

2.A Modal sensitivity analysis on continuous DMs

As discussed in Sec. 2.6, there are parts of the optical system other than the segmented mirror that can benefit from a careful sensitivity analysis performed with PASTIS. Of particular interest here are the con-

tinuous DMs in the optical path that constitute the high-order WFS&C system. Finding the natural contrast modes on these DMs could help to optimize the WFS&C strategy deployed on a high-contrast instrument. In this appendix, we look into the first explorative stages of using the PASTIS tolerancing methodology for the optimization of WFS&C techniques with continuous DMs. This was beyond the scope of the published paper but expanding PASTIS this way presents interesting possibilities for further optimizations of WFS&C, which is why we highlight some preliminary simulation results from the HiCAT testbed, and CGI on Roman. While I prepared this preliminary study on HiCAT, the intern who I supervised during my final PhD months, David Bourgeois, was investigating how to do the same thing for RST. The details about the simulators used in this section will be discussed in Chap. 3.

Below, in Sec. 2.A.1 I show how to generalize the formalism to the case of continuous DMs. Following in Sec. 2.A.2, I present the calculation of the continuous PASTIS matrix and its eigenmodes for two concrete HCI instruments. The goal here is to investigate the potential of this analysis method for a better understanding of the sensitivities and tolerances in an coronagraphic system, and to explore its prospect for modal control.

2.A.1 Formalism for continuous DMs

The mathematical formalism for the PASTIS tolerancing model has been presented in detail in Chap. 2.2. In the case in which we probe the contrast sensitivity of a continuous DM, the residual WFE on the segmented mirror are part of the phase solution for best contrast, ϕ_{DH} . Instead of an additional segmented term ϕ_s like in Eq. 2.2, we will have a phase $\phi_c(\mathbf{r})$ which is the continuous DM phase perturbation on top of ϕ_{DH} :

$$\phi = \phi_{DH} + \phi_c. \quad (2.41)$$

We have seen in Chap. 1.3.2 how the phase on a continuous DM can be described as a sum over all actuator influence functions $f_q(\mathbf{r})$ (Eq. 1.18). Assuming the coefficients a_q to be in the appropriate units, we can express the continuous DM phase ϕ_c analogously to the segmented DM phase ϕ_s in Eq. 2.1 by replacing the localized Zernikes with continuous DM influence functions:

$$\phi_c(\mathbf{r}) = \sum_{q=1}^{N_{act}} a_q f_q(\mathbf{r}). \quad (2.42)$$

With a continuous face-sheet DM, the contribution of any actuator will be highly localized but will still influence the entire DM surface. This means that the contribution of an actuator can be seen as a characteristic 2D phase map over the entire DM surface with unit amplitude, which are the influence functions $f_q(\mathbf{r})$. This influence function will be different for all the different actuators, especially when comparing actuators at the edge with actuators far away from the DM edge. In numerical simulations though, all actuators are often modeled with the same influence function, which can be approximated with a Gaussian, as is the case for the simulated DMs on HiCAT, and the pupil-plane DM used in the optical simulations of CGI on RST, examples of both of which are shown in Sec. 2.A.2.

Developing the DM phase ϕ_c in the small aberration regime allows us to follow the same development of the PASTIS forward model like shown in Eqs. 2.3 – 2.8, leading to the PASTIS formula to evaluate the contrast as a function of pupil-plane aberrations in Eq. 2.9. The difference in this case is that \mathbf{a} represents the aberration vector of all the coefficients a_q for the continuous DM influence functions in Eq. 2.42. The PASTIS matrix M now describes the differential contrast sensitivity per unit aberration of each pair of continuous DM actuators, and it needs to be calibrated with measurements of the contrast response of actuator motions, analogously to the segmented case described in Chap. 2.2.2. In the presence of a best-contrast DH solution ϕ_{DH} , the matrix elements m_{ij} that are explicitly shown in Eq. 4.5 for the segmented case, now take the form of:

$$m_{ij} = \langle C\{P(\mathbf{r}) e^{i\phi_{DH}(\mathbf{r})} f_i(\mathbf{r} - \mathbf{r}_i)\} C\{P(\mathbf{r}) e^{i\phi_{DH}(\mathbf{r})} f_j(\mathbf{r} - \mathbf{r}_j)\}^* \rangle_{DH}. \quad (2.43)$$

As laid out in Chap. 2.2.2, the PASTIS matrix can thus be calculated from individual average DH contrast values from intensity images, obtained by aberrating all pairs of segments, or in this case actuators, one by one.




Figure 2.19: Final DM surface commands for a baseline contrast $c_0 = 6 \times 10^{-8}$ obtained with iterative pair-wise estimation and stroke minimization on a non-segmented CLC simulation of the HiCAT testbed without any WFE. The DM surfaces are the main contributor to the phase term ϕ_{DH} in Eq. 2.43, describing the elements of the PASTIS matrix. The pair-aberrations are applied on top of these surface commands when the PASTIS matrix is built and the resulting average contrast is measured in the circular DH with an extent of $6\text{--}11 \lambda/D_{LS}$.

2.A.2 Simulated results on HiCAT and RST

The HiCAT simulation setup used for the preliminary work presented in this appendix is described in more detail in Chap. 3.3.1, but without the segmented DM in the beam, thus operating with a *non-segmented* classical Lyot coronagraph (CLC), with a simple circular pupil. The simulations include no WFE and are run at a monochromatic wavelength of 640 nm. Similarly to the WFS&C procedure used in Chaps. 3.3.1 and 4.3.2, a pair-wise estimation and stroke minimization loop is run first to obtain a baseline contrast, setting $c_0 = 6 \times 10^{-8}$, using two-DM control. The resulting coronagraphic image and the DM commands from the final iteration are shown in Fig. 2.19. These DM commands are loaded each time prior to performing a PASTIS analysis on the HiCAT simulator, and they are the dominating part of the phase term ϕ_{DH} .

The results presented here show the sensitivity analysis performed on the in-pupil continuous DM made out of 952 individual actuators, although many of the actuators located at the edge of the pupil remain hidden by various pupil masks. Contrary to the numbering of the segments on the segmented DM, the continuous DM actuators are numbered row by row, which needs to be kept in mind when interpreting the results visually. Following the development in Chap. 2.2.2, we push each non-repeating actuator pair (including hidden actuators) with a calibration aberration of $a_c = 10$ nm, recording $n_{meas} = 453, 628$ intensity measurements, from which we construct the PASTIS matrix shown in Fig. 2.20, left. As in previous examples, the matrix diagonal is very pronounced, describing the contrast contribution from each individual actuator alone. The disconnected nature of the diagonal comes from the choice of actuator numbering: since the actuators are numbered from left to right and bottom to top, at the beginning and end of each row we are poking actuators that lie outside of the illuminated HiCAT pupil, causing little to no effect on the DH contrast. As this repeats periodically, this shows as interruptions of contrast influence as we work our way through all actuators. The inversion of the matrix, like shown in Chap. 2.3, yields the eigenvalues shown in Fig. 2.20, right.

The eigendecomposition of the PASTIS matrix (Eq. 2.18) results in orthonormal eigenmodes with associated eigenvalues. We can see in the plot of eigenvalues in Fig. 2.20, right, that these relative sensitivities span a large range of seven orders of magnitude, split in four regimes, indicated by the distinct bumps in the figure. A subset of modes from the first of these, with the highest eigenvalues, are displayed in Fig. 2.21. As discussed in Sec. 2.3.1, the eigenvalues represent the coronagraph sensitivity to the mode, so the modes shown in Fig. 2.21 have a high impact on the contrast, relatively speaking. They all display highly symmetric features. This high-impact regime seems to be contained to the first bump in Fig. 2.20, right, its limit being marked by the vertical dashed line. Any modes beyond this point show a high level of noise, as can be seen in the bottom right mode in Fig. 2.21, which lies beyond the high-impact regime, at index 338.

We compare these results to simulations from CGI on RST, which were done using the WebbPSF package (see also Chap. 3). Using a Shaped Pupil Coronagraph (SPC) and assuming a WFE-free system, the




Figure 2.20: HiCAT PASTIS matrix (left) and eigenvalues (right) for the pupil-plane continuous DM with 952 individual actuators. *Left:* The chopped-up structure in the matrix is a result of the actuator numbering on the DM: instead of going outwards in spirals starting from the center, the actuators are numbered row by row. Since the actuators at either end of each row are not illuminated, their influence on contrast is negligible. *Right:* The eigenvalues span almost seven orders of magnitude and suggest that there are at least three different eigenmode regimes with different relevance for the contrast. The dashed vertical line marks the point where noise starts dominating the eigenmodes.




Figure 2.21: Pupil-plane eigenmodes for the continuous PASTIS analysis on HiCAT, with their respective mode index corresponding to the eigenvalues in Fig. 2.20, right. Modes with an index up to ~ 270 , marked by the vertical dashed line in Fig. 2.20, have a significant impact on the coronagraph contrast, as quantified by the respective eigenvalue.




Figure 2.22: Roman CGI simulation setup with a Shaped Pupil Coronagraph (SPC), shown on the far left, as projected into the exit pupil of RST. The middle and right plots show the coronagraphic contrast without and with a DH mask applied, respectively. Both plots are shown on the same scale. This setup provides a baseline contrast of $c_0 = 5 \times 10^{-9}$ in a symmetrical wedge-shaped DH from $3-9\lambda/D$, under the absence of any WFE in the simulator.

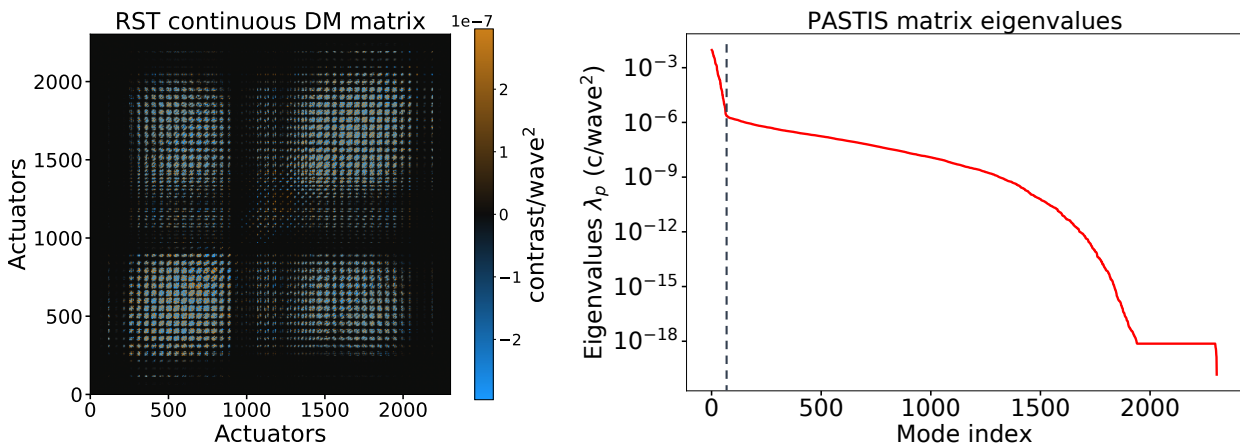


Figure 2.23: *Left:* PASTIS matrix for the continuous CGI analysis on RST, using the SPC pictured in Fig. 2.22, on a highly saturated scale in order to show the faint features. *Right:* Eigenvalues of the PASTIS matrix shown to the left. The dashed vertical line marks the point where noise starts dominating the eigenmodes. Credit: David Bourgeois.

coronagraphic pupil mask and the resulting coronagraphic PSFs are shown in Fig. 2.22, yielding an average DH contrast of $c_0 = 5 \times 10^{-9}$ at a wavelength of 770 nm. As opposed to the HiCAT simulations further above, where the high-contrast regime is attained by deploying an iterative WFS&C algorithm, the RST setting used here creates the DH instantaneously with the coronagraph alone. With a 48×48 actuator continuous DM, we calculate the PASTIS matrix shown in Fig. 2.23, left, with a calibration aberration of $a_c = 5$ nm, and show its eigenvalues in the same figure on the right. In the PASTIS matrix for the SPC on Roman, we can see distinct regions of actuator pairs that have the most pronounced impact on the coronagraphic contrast. These matrix regions correspond to pairs in which either one or both of the poked actuators lie in the two more transparent areas of the shaped pupil mask as shown in Fig. 2.22, left. The rest of the actuators is very concealed by the pupil optics and have thus less influence on the contrast. In the same way like for HiCAT, the DM actuators on Roman CGI are numbered row by row over the entire pupil, which is what causes the choppy visual representation of the PASTIS matrix. The eigenvalues of the RST matrix span twice as many orders of magnitude compared to HiCAT, but they also display much more pronounced regimes. Up until eigenmode with index 80, marked by the vertical dashed line in Fig. 2.20, right, the eigenmodes show distinct features and the coronagraph is more sensitive to them than higher-index modes. A sample of the modes from this regime is shown in Fig. 2.24. The shown modes are particularly pronounced in parts of the pupil that remain unconcealed.




Figure 2.24: Pupil-plane eigenmodes for the continuous PASTIS analysis on RST with an SPC, with their respective mode index corresponding to the eigenvalues in Fig. 2.23, right. Credit: David Bourgeois.

3

SIMULATING TOLERANCING ANALYSES OF SEGMENTED APERTURE CORONAGRAPHS

The output and results of my PhD work were enabled by the software I developed. In this chapter, made of three sections, I outline the details on numerical simulation studies I conducted in the context of coronagraph tolerancing, and the tools I developed for this purpose. In Sec. 3.1, I describe the software package I wrote for optical simulations and coronagraph tolerancing, in Sec. 3.2 I present an early tolerancing analysis of a JWST coronagraph, and in Sec. 3.3 I show the results of simulating contrast sensitivity experiments on the HiCAT testbed.

For the initial development of the semi-analytic PASTIS matrix, I used the open-source Python package WebbPSF to simulate the effect of segmented aberrations through a NIRCcam coronagraph on JWST. However, JWST coronagraphy is not designed to reach the high-contrast regime needed for exoEarth detection (10^{-10}). Instead, I continued my work by creating a numerical model of the LUVOIR A telescope. The objective seemed straightforward: take a coronagraph simulation and add segmented aberrations to it. However, this proved non-trivial in practice. The only open-source segmented DM model available at the time in Python was within the POPPY optical propagation library, but it was incompatible with the APLC simulations for the LUVOIR coronagraph design from the segmented coronagraph design analysis (SCDA) study. Thus, I started developing my own segmented DM model with special focus on compatibility with pupil plane coronagraph optics. This led to the creation of a custom Python package for the full process of a coronagraph sensitivity analysis, the PASTIS package, which I talk about in the first part of this chapter. It has been published through Zenodo and regular releases appear as the tool is being continuously improved (Laginja 2021).

After concluding the development of the tolerancing tool, I went back to apply it to NIRCcam, the results of which I show in the second part of this chapter. While my simulations include somewhat simplified assumptions, the order-of-magnitude accordance with the design and ground-testing results of a real space telescope are encouraging for more in-depth work.

The third and last part of this chapter shows the simulation results I obtained while preparing for the experimental validation of the PASTIS tolerancing model on HiCAT, which I presented as a poster and in conference proceedings during the SPIE Astronomical Telescopes & Instrumentation conference in December 2020, remotely during the Coronavirus pandemic (Laginja et al. 2020). Writing the code for the experiments on the HiCAT testbed emulator meant that I was directly preparing the hardware operations, since HiCAT has a dual-mode operation setup which uses the same scripts in simulation and on hardware. These simulations revealed the need to find a way to work with a differential contrast contribution, stemming from the isolated effect of the segmented DM, rather than the total contrast that will be influenced by other WFE sources as well, such as environmental drifts. This problem was later addressed in Chap. 4.

With the tolerancing analysis, the LUVOIR A simulator and the interfaces for the JWST and HiCAT simulators, the PASTIS analysis package was 100% my work, from the implementation, over the integration of unit and regression tests, to the overall software management. Laurent Pueyo then contributed simulator models for LUVOIR B and HabEx, both with vortex coronagraphs, and many other

improvements followed. The PASTIS package is a central piece of the work of our team's industry collaboration for LUVOIR tolerancing (Pueyo et al. 2021; Sahoo et al. 2021; Sahoo et al. 2022), and it has already been used externally in published studies (Pogorelyuk et al. 2021). I intend to leverage these tools for my work after my PhD and to continue developing them.

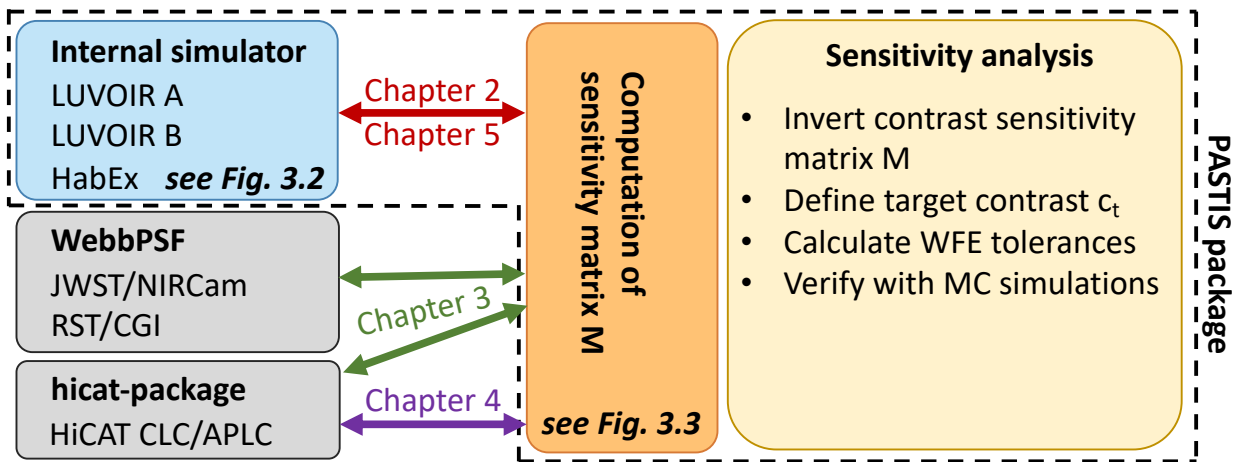


Figure 3.1: High-level functionality of the PASTIS software package, which I developed from scratch during my PhD. It performs a sensitivity analysis from a coronagraphic instrument simulator. It can either use internally scripted telescope simulators, or interface with third-party optical propagators like WebbPSF or hicat-package. The colored arrows indicate which simulator was used in which chapter of this thesis.

3.1 A numerical tool for statistical WFE tolerancing

From initially only providing the LUVOIR A telescope architecture with its three distinct APLC designs that was used in Chap. 2 to demonstrate the segmented tolerancing analysis with PASTIS, the “PASTIS” package has grown into a much larger infrastructure for end-to-end modeling of coronagraphs and the impact of aberrations. This package enables a numerical study of coronagraphic sensitivity across various different telescope concepts and aberration sources; a schematic of it is shown in Fig. 3.1. The goal of the PASTIS Python package is to provide tools with which one can perform a full tolerancing analysis on a specific modal basis. The three main components allowing us to do that are:

1. numerical propagators that correctly integrate a telescope aperture with a coronagraph and an aberration mode basis of choice (left-hand side in Fig. 3.1),
2. the capability to calculate contrast sensitivity (PASTIS) matrices for aberrations with various spatial frequencies (center block in Fig. 3.1),
3. the calculation of statistical tolerances on aberrations for an input target contrast, together with verification methods like Monte Carlo simulations (right-hand side in Fig. 3.1).

The separation of the numerical simulators from the full analysis of PASTIS matrices enables an easy swap-in of third-party optical simulators to create PASTIS matrices. These can be directly analyzed with the tolerancing model, which I exploited for contrast tolerancing analyses of a JWST coronagraph, and the segmented HiCAT coronagraph. I show some results from the tolerancing of a coronagraph on NIRCcam in Sec. 3.2. The JWST simulations I used were intentionally simplified for proof-of-concept simulations. As I was moving on to more realistic setups, I leveraged the work I invested in the optical simulator and the control interface of HiCAT in order to predict the performance of a segmented coronagraph in a HiCAT-like environment, shown in Sec. 3.3, before performing the experimental validations presented in Chap. 4.

The overall simulation package shown in Fig. 3.1 is a generalized simulation tool for the analysis of coronagraphic sensitivities. It implements a streamlined interface for the computation of PASTIS matrices and their inversion, followed by the calculation of tolerancing results as a function of target contrast, finalized by cross-checks and validations with Monte Carlo (MC) simulations. While the analysis procedure can interface with any optical simulator, the PASTIS package itself contains a full numerical simulator for some specific telescope designs like the segmented LUVOIR A and B, with APLCs, and the monolithic HabEx with a DM-supported vortex coronagraph (DMVC), together with various deformable mirrors that work with different aberration mode bases (see Sec. 3.1.1). These optical simulators provide the capability

to model various aperture–coronagraph combinations, with different sets of optical aberrations. The initial goal of the code was to simulate segment mode misalignments with pistons on the LUVOIR A telescope, as this was the use-case in the formal development of the theoretical framework. However, the package now supports a range of local and global modes, the latter of which are also implemented on monolithic apertures (HabEx, RST) for extensive sensitivity studies (see also Chap. 5.2).

3.1.1 Implemented telescopes, coronagraphs and aberration modes

The optical propagation of different coronagraphs, including the APLC, has been implemented by various other simulators, as have been models of segmented mirrors. The problem that the end-to-end simulator in the PASTIS package solves is that of proper sampling of the (aberrated) segments themselves in the telescope pupil with respect to coronagraph optics located in a pupil plane as well. Having apodizers, Lyot stops and DM phase maps used for example in APLCs and DM-supported VCs means that all these pupil-plane optics have to be aligned correctly to each other, as well as to the individual segments of the telescope pupil, otherwise the coronagraph performance will be degraded and the propagation of WFE incorrect. For coronagraphic propagations under the absence of segmented aberrations, it suffices to get the pupil sampling right globally, meaning that all pupil-plane optics of a coronagraph do not exhibit lateral misalignments or distortions. Considering that the coronagraph optics are usually designed as a unit, this is easily achievable in practice. For propagations that include segment-level aberrations however, which is the core topic regarded by the PASTIS model, it is also necessary to ensure the correct sampling of each individual segment within the pupil. Each individual pixel of the segmented aberration basis has to align with each individual pixel of the segments in the telescope aperture, as well as any of the other pupil-plane optics (apodizer, LS, DM commands). In this way, mid-spatial frequency aberrations stemming from the aperture segmentation propagate correctly through the coronagraph.

The PASTIS package, and its end-to-end simulators, are fully written in Python, and use the HCIPy package (Por et al. 2018) for optical propagations (Fraunhofer and Fresnel), including the Lyot and vortex coronagraph implementations. While other optical simulation packages exist, HCIPy was found suitable due to its flexibility in optical and physical units (e.g., λ/D units in the focal plane, SI units in the pupil plane), rather than being tied to instrumental parameters of a full optical system (e.g., arcsec in the focal plane, which requires a plate scale). Further, HCIPy is coded very modular, which makes it possible to use specific parts of it without having to surrender to a code basis that requires the use of exclusively its own features. Object-oriented programming (OOP) is being used in the PASTIS package to create a class structure that stays modular and compact. In particular within the optical simulator modules of the package, this allows easy inheritance of the most important features of any given aperture and active optic (e.g., segmented or continuous DMs). The class inheritance diagram of the available optical simulators within the PASTIS package is shown in Fig. 3.2. The main base class is `Telescope`; it sets up all global components like the total pupil diameter, operating wavelength and image sampling. This class also contains an aberration basis of global Zernike modes (Noll 1976), numerically implemented to operate as a DM whose influence functions are global Zernikes. In the same way, this class contains a DM implementation with global ripples (sine wave aberrations) as influence functions, an actual continuous DM with individual actuators with a variable actuator count covering the entire pupil, and an out-of-band Zernike WFS (OBWFS).

All following classes inherit all attributes and methods from `Telescope` and adapt them to their respective implementation. The class `SegmentedTelescope` adds attributes that define the respective segment positions, sizes and geometries, which is the basis on which it implements the two optional segmented DMs. The segmented DM that is instantiated by default (“SegDMptt”) provides piston, tip and tilt (PTT) control of each individual segment, and requires less computation time for a single PSF calculation than the second segmented DM (“SegDMmulti”). The user can choose to manually instantiate the SegDM-multi option with an arbitrary number of local Zernike modes on each segment. This provides enhanced capabilities for the segmented DM, but also turns all optical propagations much more computationally intensive compared to usages with the PTT-only segmented DM. Similarly, the user can choose to add a segmented DM with a custom local mode basis obtained from L3 Harris Technologies (“HarrisSegDM”, see Chap. 5.2.3).

The implementation of concrete telescope architectures and coronagraph designs is done via inheritance from these two base classes. To make implementations of APLCs easier on segmented apertures, the package contains an abstract class `SegmentedAPLC` that inherits from `SegmentedTelescope`, which handles an apodizer, FPM, LS, IWA, OWA and DH mask. One of the three currently implemented telescopes is LUVUOIR A with the class `LuvOirA_APLC`, which pipes the correct apodizer, LS file and coronagraph parameters through to `SegmentedAPLC`. It also provides the option to choose between the three APLC designs (small, medium and large) with a single class parameter. The implementation of `LuvOirBVortex` exploits the infrastructure of `SegmentedTelescope` to correctly implement the segment actuations with a Vortex coronagraph. Finally, since HabEx has a monolithic pupil, the class `Habex_VVC` inherits directly from `Telescope` to implement its DM-supported VC on a circular aperture.

The source files (aperture, apodizers, Lyot stop) of the LUVUOIR A APLC suite are credited to the Segmented Aperture Coronagraph Design and Analysis (SCDA¹) study, while the LUVUOIR B and HabEx data were provided by Garreth Ruane, and functionally merged into the PASTIS package by Laurent Pueyo. For each telescope instance, the simulator calculates the electric field at each pupil and focal plane of the optical system and returns them to the user. Each of the DMs (global Zernike, ripple DM, continuous DM, `SegDMptt`, `SegDMmulti`, `HarrisSegDM`) can be controlled separately, and they are folded into the optical propagation process, which also has the option to be normalized to a single photon in the entrance pupil. The generic `SegmentedTelescope` can also be used alone, without a coronagraph.

Some points that show potential for improvement of this simulator include creating an abstract class `SegmentedVC`, analogous to `SegmentedAPLC`, which can serve as a base class for any segmented telescope with a Vortex coronagraph. Further, a general refactor of the code could see an abstraction of the telescope structure and coronagraph, making them separate. In such a case, a realistic telescope implementation could use multiple inheritance to mix-and-match a telescope primary mirror (monolithic or segmented) with a coronagraph type (e.g., APLC, VC, ...).

3.1.2 Statistical sensitivity analysis

For the calculation of PASTIS sensitivity matrices, a simple class inheritance structure starting with the abstract class `PastisMatrix` allows the computation of matrices for distinct telescope simulators through subclassing, as shown in Fig. 3.3. This allows the package to support the calculation of contrast sensitivity matrices for the LUVUOIR and HabEx simulators described in Sec. 3.1.1. Furthermore, some additional telescope/coronagraph combinations can be handled through an interface with external simulators. This includes NIRCcam coronagraphs for JWST and Roman CGI with the open-source package `WebbPSF`, and the HiCAT simulator which is a private package of the Makidon Lab team. Ongoing work is aiming to make the integration of new simulators even smoother through a more unified interface between such simulators and the matrix calculation classes.

The third and final part of the package concerns the statistical tolerancing analysis which is performed with a PASTIS matrix and a target contrast as inputs. Since these calculations require only linear algebra, no specific adjustments per individual telescope simulator is needed. Only the optional visualisation of the results, and sanity checks with Monte Carlo simulations require an integration of the telescope simulator in the analysis section as well.

In the following, some results are presented that use the code interface provided by the PASTIS package. Sec. 3.2 shows the sensitivity analysis of a coronagraph on JWST, operating on less ambitious contrast levels than envisioned for RST or LUVUOIR. Finally, Sec. 3.3 presents simulations with the HiCAT simulator, in preparation for the experimental validations in Chap. 4.

3.2 Contrast sensitivities to segment aberrations on a JWST coronagraph

Early simulations for the pair-aberrated contrast influence analysis with PASTIS were performed by generating analytical images (Lebouilleux et al. 2018b). These were constructed from a combination of the

¹<https://exoplanets.nasa.gov/exep/technology/SCDA/>

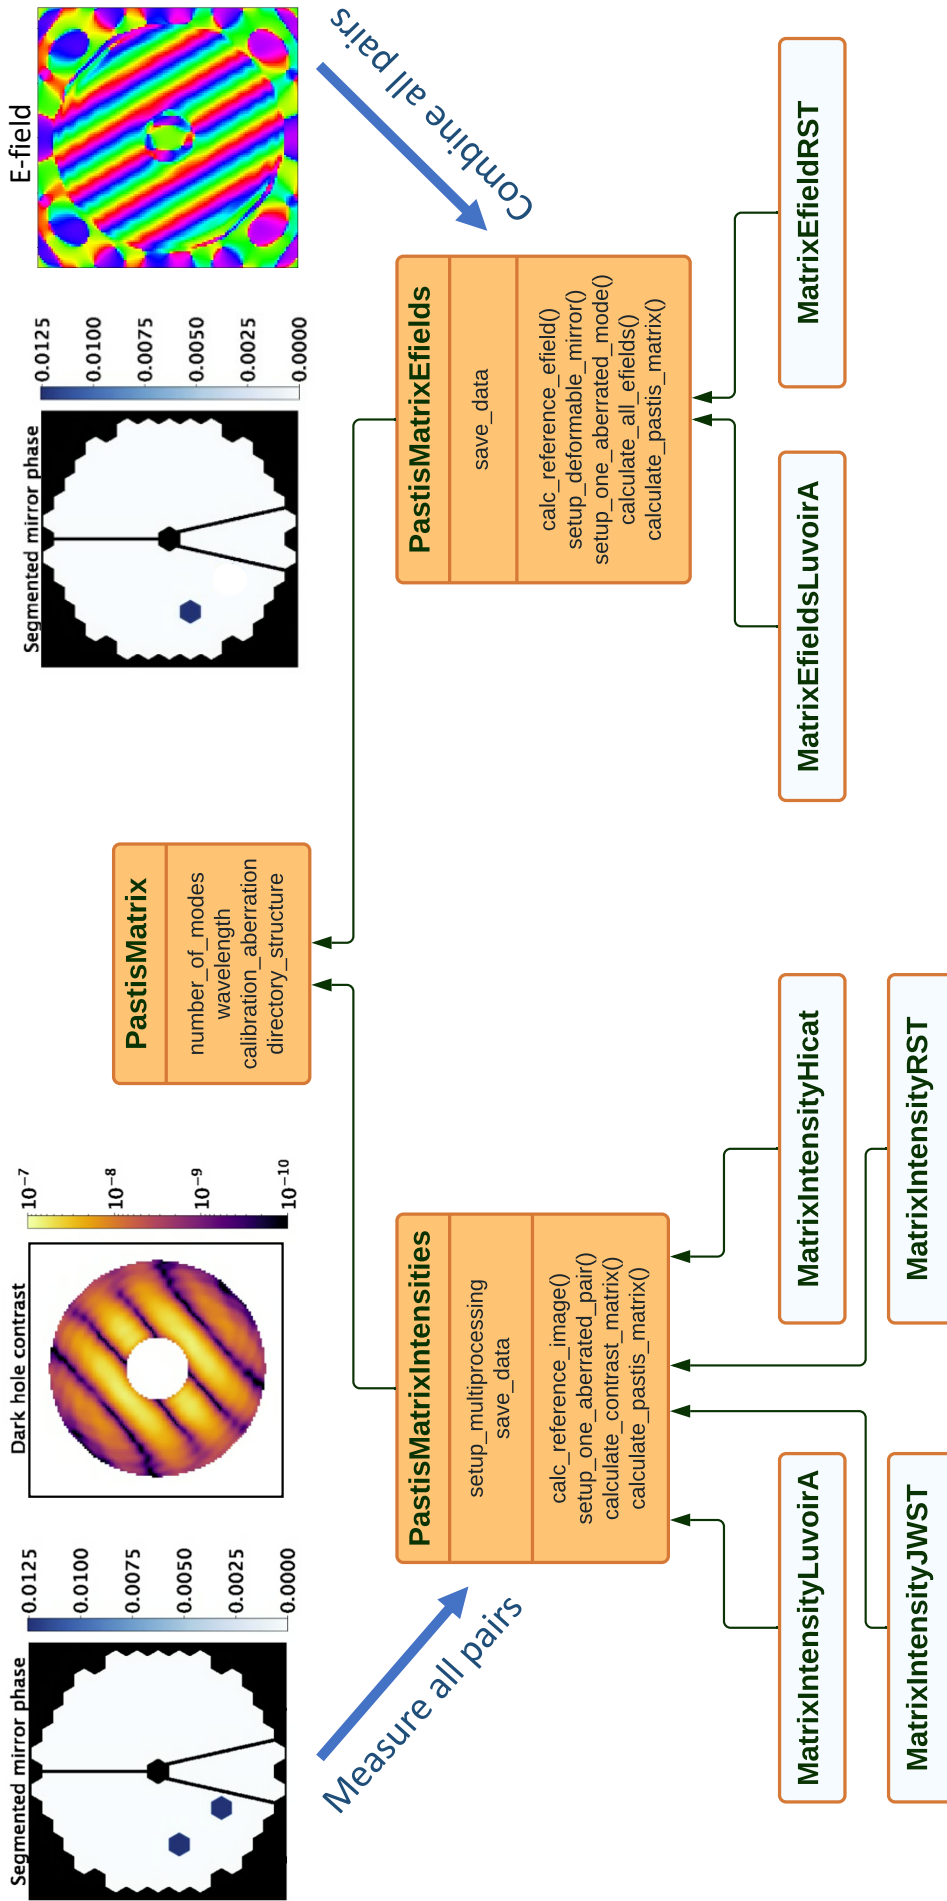


Figure 3.3: Class inheritance diagram for the computation of PASTIS matrices. The abstract parent class `PastisMatrix` handles the basic optical parameters that have no dependence on the specific simulator used. Two abstract child classes implement the calculation of PASTIS matrices either through measurements of focal-plane intensities from pair-aberrated segments/modes, or through measurements of focal-plane E-fields from individually aberrated segments/modes. Measuring PASTIS matrices from pair-aberrated segments/modes is more reliable on hardware, while combining them from individual E-fields is more efficient in simulation. Concrete child classes for each telescope allow the calculation of a full PASTIS matrix.

single-segment intensity and the respective cross-terms arising from two segments being pushed at the same time. This was done on a proxy aperture for future large space telescopes, a 36-segment aperture without a center segment. To derive sensitivities for an existing telescope and coronagraph, and prepare the tolerancing model for LUVVOIR, I started an implementation of the analysis technique on JWST with the open-source simulator WebbPSF (Perrin et al. 2014a). Since the fully analytical image models used in Leboulleux et al. (2018b) required a calibration with an end-to-end (E2E) simulator in order to assemble the final PASTIS matrix (see Leboulleux et al. 2018b, Sec. 4), I extended the model from a purely analytical to a semi-analytical one, in which the E2E simulator is used for the image generation directly, from which the PASTIS matrix is then calculated analytically (Laginja et al. 2019). The interest of this work does not lie primarily in the derivation of realistic segment co-phasing tolerances for JWST, although it allows us to make some qualitative conclusions about the sensitivity range of the NIRC*am* coronagraphs. Considering its anticipated coronagraph performance and overall WFE, JWST is expected to observe at the 10^{-5} raw contrast level (Perrin et al. 2018), far away from the high-contrast regime needed for exoEarth detection. However, the considerations that went into the conception of a large segmented telescope that can support coronagraphy, and the simulation tools of such, are valuable lessons in the preparation for other, more performant missions.

3.2.1 Coronagraphy with NIRC*am*

WebbPSF provides an optical model of all JWST instruments, including their coronagraphs. JWST is an 18-segment observatory without a central pupil segment, and coronagraphy is supported both for the MIRI and NIRC*am* instruments. Both of the two identical NIRC*am* modules (A and B) contain the same set of three round, and two bar occulters, which get paired with two separate Lyot stops per occulter type to form a classical Lyot coronagraph (Krist et al. 2010, 2007). NIRC*am* coronagraphy covers a wavelength range from 2.5-5 μm , which overlaps in the shorter wavelengths with ground-based facilities like SPHERE and GPI. While JWST will have a much higher absolute sensitivity, these ground-based imaging instruments will continue to have higher angular resolution and better inner working angles, which puts the science interest with the JWST coronagraphs more on the middle and longer wavelengths available to NIRC*am*. The round occulter dedicated to a wavelength of 3 microns has the most science filters available, so we use it to perform a contrast sensitivity analysis with respect to segment co-phasing errors on JWST. It is used with the round-occulter Lyot stop, operated at a central wavelength of 3.35 μm , and with an inner working angle of $6\lambda/D$, which corresponds to 0.63 arcsec on-sky.

Since the objective of this analysis is to identify the influence of segment actuation on the coronagraphic image contrast on NIRC*am*, we eliminated any other aberration sources by taking out any science instrument internal WFE in the simulator, and we assumed a perfectly phased starting position of the segments. In such a scenario with the CLC on NIRC*am*, we use the round occulter identified by “MASK335R” with the Lyot stop for these round occulters, “CIRCLYOT” (see Fig. 3.4), with a filter centered at 3.35 μm called “F335M” to obtain coronagraphic images. We define the DH in these images, shown in Fig. 3.4, from an IWA of $6\lambda/D$ to an arbitrarily defined OWA of $16\lambda/D$ (within the NIRC*am* field of view). Inside this focal-plane region, the described setup yields an average coronagraph floor of $c_0 = 2.5 \times 10^{-7}$.

3.2.2 Sensitivity analysis on NIRC*am* with a CLC

Proceeding in the same way like described in Chap. 2, we sequentially aberrate all non-repeating segment pairs and measure the resulting average contrast in the coronagraphic DH. We limited the analysis to the contrast influence of pure local piston errors on the segments. The calibration aberration had to be chosen above the minimal actuator move on the JWST segments, which is 5 nm, and also high enough to create a significant contrast change in the DH beyond the contrast floor. This led to the calibration aberration a_c being set to 100 nm of WFE per segment, which corresponds to a global WFE over the pupil of 22.5 nm rms for a single pistoned segment, and to 33.8 nm rms for a pair of pistoned segments. After measuring the contrast response of all aberrated pairs, we calculated the PASTIS matrix (Fig. 3.5, left) and performed a validation by evaluating global pupil aberrations caused by segmented pistons, scaled to a range of different global rms values, which yields a hockey-stick-like curve (Fig. 3.5, right). Like the PASTIS




Figure 3.4: Simulations of the NIRCcam circular-occulter coronagraph at $3.35 \mu\text{m}$. *Left:* Lyot stop with a $\sim 20\%$ geometrical transmission. *Center:* Diffraction in the Lyot plane. *Right:* Non-coronagraphic (top) and coronagraphic (bottom) PSFs in absence of any wavefront errors. We note that they are not displayed on the same contrast scale. The area limited by the white dashed rings indicates the DH used in the PASTIS analysis. The average contrast over this area is $c_0 = 2.5 \times 10^{-7}$.




Figure 3.5: *Left:* PASTIS matrix for NIRCcam instrument on JWST, with round occulter coronagraph “MASK335R”, calculated with the PASTIS Python package. *Right:* Validation of the semi-analytic (SA) PASTIS matrix to the left by comparing results of Eq. 2.36 with numerical propagations from WebbPSF. The resulting hockey stick-shaped curves overlap up to a global WFE rms value of 1×10^{-1} waves.




Figure 3.6: Eigenmodes of the NIRCcam PASTIS matrix, ordered from highest (top left) to lowest (bottom right) influence on the coronagraphic contrast with a round-occluder CLC.

matrices for segmented telescopes shown previously, the NIRCcam matrix features a pronounced matrix diagonal that indicates the contrast influence of single aberrated segments, with the off-axis elements denoting the cross-terms of the intensity influence between two segments. Adjacent segments aberrated at the same time influence the contrast less than when the segments are spatially separated, leading to the negative blue off-axis streaks in the matrix. The hockey stick curve shows that the contrast, calculated with Eq. 2.9 from an arbitrary segmented aberration vector, yields the same results as measuring the contrast after propagating the same aberration vector with the E2E simulator, for aberrations smaller than 1×10^{-1} waves. The flattening of the E2E curve beyond this point is caused by the phase wrapping on individual segments, an optical effect that appears in the E2E simulator, but not in the analytical PASTIS equation.

The eigendecomposition of the matrix in Fig. 3.5 results in 18 optical influence modes shown in Fig. 3.6, ordered from highest to lowest influence on the DH contrast. The eigenmodes of the NIRCcam PASTIS matrix display very similar characteristics to the ones we observed with the LUVUOIR coronagraph in Chap. 2.3.1, in particular with respect to the segments that are fully or partially obstructed by the LS or support struts. The modes with the least influence on the contrast, on the bottom of the figure, appear to be approximated low-order Zernikes with the more concealed segments particularly pronounced. The most influential modes on top of the figure bring forward the segments that are the least obscured by the pupil-plane features of the telescope and coronagraph.

With the modes being orthonormal, we can demonstrate their additive contribution to the DH contrast with Eq. 2.26. This is shown in Fig. 3.7, left, after rescaling each mode weight with the inverse of their respective eigenvalues, which scales them in such a way that each mode contributes equally to the total contrast. This contrast allocation is chosen arbitrarily, and other choices can be made when pursuing a concrete error budget. Following the development in Chap. 2.5.2, instead of allocating uniform contrast to each eigenmode, we can choose to allocate a uniform contrast across all (independent) segments, which transfers to a non-uniform contrast allocation across the modes. However, the modes remain additive in contrast, and a comparison between these two contrast allocation strategies, as projected into the basis of the eigenmodes, is shown in Fig. 3.7, right.

Assuming independent segments makes it possible to calculate standard deviation limits on each segment in the pupil for a particular target contrast to be maintained as the mean over many states of the segmented mirror. Using Eq. 2.36, we use the PASTIS matrix to calculate the per-segment tolerances for a target contrast of $c_t = 10^{-6}$ and plot the results in Fig. 3.8, left and middle. The individual segment WFE tolerances for a total mean contrast of 10^{-6} range from 28 nm to 129 nm. We can see directly how there is a differentiation between the segments when it comes to their respective tolerances: the most relaxed

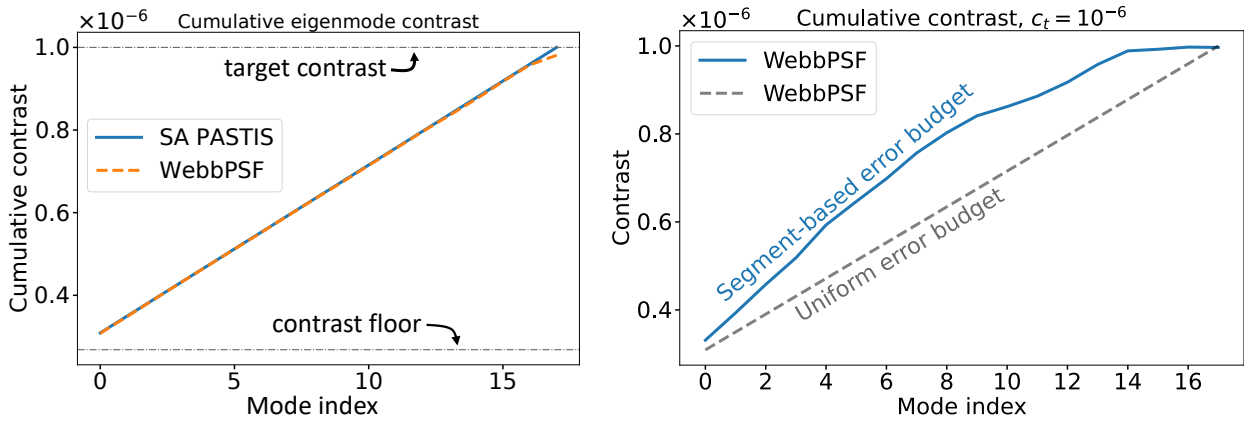


Figure 3.7: *Left:* Uniform contrast allocation across all eigenmodes, as calculated with the PASTIS Eq. 2.36 with the JWST PASTIS matrix shown in Fig. 3.5 (solid blue), and with the WebbPSF simulator (dashed orange). *Right:* Comparing the uniform contrast allocation across all eigenmodes (dashed grey) from the left side with a segment-weighted allocation to all eigenmodes (solid blue), in which all independent segments in the pupil contribute equally to the final contrast. Both curves were computed with the WebbPSF simulator.




Figure 3.8: Segment requirements, expressed as standard deviations, and MC simulation histogram for JWST, as a line plot (left) and as a map with labelled segment numbers (center). The most restricted WFE requirements for a target contrast of 10^{-6} are 28 nm on segments 2 and 6, which have the highest LS throughput (see Fig. 3.4). The most relaxed segments are numbers 7, 12 and 14 who have a WFE requirement of ~ 90 –130 nm and are completely concealed by the LS, and also partially hidden behind secondary mirror support struts. The MC simulation with 1000 independent WFE maps as prescribed by the tolerancing analysis recovers a mean contrast of 10^{-6} (right).

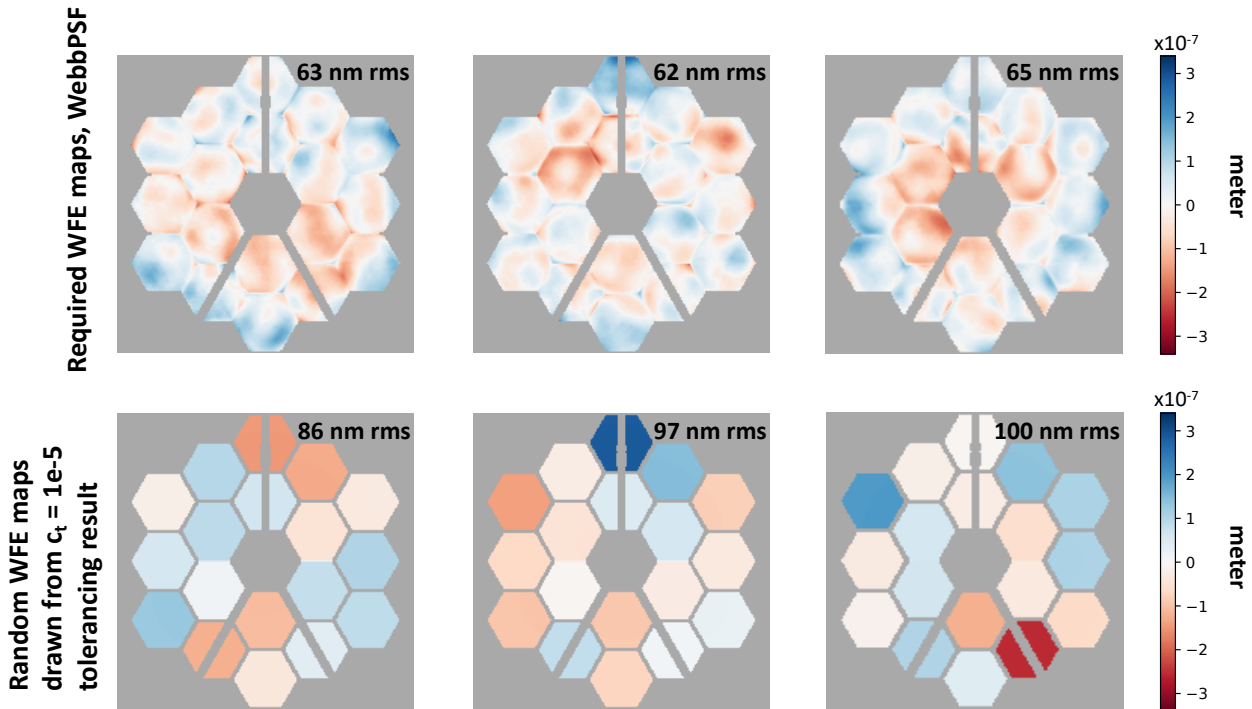


Figure 3.9: Direct comparison between three arbitrary required WFE maps from WebbPSF (top), and three random WFE maps as generated for the 10^{-5} MC simulation (bottom), with their respective global WFE rms values indicated in each upper right corner. The WFE maps from the tolerancing simulations presented here (bottom row) assume that all WFE of the system lies in segment-level piston. The mean contrast from 1000 such maps is 10^{-5} , which means that the requirements taken from WebbPSF (top row), which are more restrictive, are sufficient to produce that target raw contrast. We can observe on the exemplary maps on the bottom that certain segments get away with a larger WFE, for example segments behind support struts, or in the outer ring of the pupil.

segments are those that are both fully concealed by the Lyot stop and are additionally also covered by the secondary support struts, segments 7, 12 and 14. The most restricted segments with a WFE requirement of 28 nm are segments 2 and 6, which see the highest transmission through the very particularly shaped LS, see also Fig. 3.4. Using these tolerancing limits in a Monte Carlo simulation with a set of 1000 segment piston realizations, we plot the histogram shown in Fig. 3.8, right. Analogously to the MC experiment in Chap. 2.5.1, we use the segment standard deviations to draw 1000 normally distributed aberration maps, and use the E2E simulator to record the resulting average contrast in the DH. The mean value of all these propagations is 1.01×10^{-6} , which meets the target contrast of $c_t = 10^{-6}$ to less than a 1% error, confirming the standard deviation tolerances of the individual segments.

While these simulations were simplified by omitting any WFE sources other than from the segmented Optical Telescope Element (OTE), they give a sense of the order of magnitude for the sensitivity range of the NIRCcam coronagraph to segment misalignments. Since WebbPSF contains model examples of WFE maps for the OTE², we can compare the results from the PASTIS tolerancing analysis to the real error budget of JWST. The modeled OTE WFE in WebbPSF consist of ten instances of probabilistic maps each for the predicted, and for the required WFE performance. The median of the total WFE rms of the ten predicted performance maps is 70 nm rms, with the required maps having a more conservative median WFE rms of 90 nm rms. Three arbitrary examples of the required WFE maps in WebbPSF are shown in the top row of Fig. 3.9. We can compare this directly to some of the randomly created WFE maps from a MC simulation like that shown in Fig. 3.8, right, drawn from the requirement map displayed in the middle of the same figure. The median global WFE rms value of all 1000 WFE map realizations is 52 nm rms, defining the WFE requirement of our analysis. This lies below the predicted, and well below the required OTE WFE for NIRCcam, suggesting that this instrument will not be able to reach the requested raw contrast

²<https://webbpsf.readthedocs.io/en/latest/jwst.html#optical-telescope-element-ote>

level of 10^{-6} . However, this performance level is distinctly better from what is expected of NIRCcam at $\sim 10^{-5}$ (Perrin et al. 2018). Repeating the tolerancing analysis for $c_t = 10^{-5}$ (calculate individual segment tolerances from the PASTIS matrix in Fig. 3.5, perform MC simulation with 1000 randomly drawn maps from these tolerances), the median of these 1000 random WFE maps now amounts to 187 nm rms, which is above the numbers of the WebbPSF-internal WFE maps from both predicted as well as required modeling of the OTE. Three examples of these random WFE for $c_t = 10^{-5}$ are displayed in the bottom row of Fig. 3.9. We can thus conclude that the OTE on JWST will perform better than our tolerancing analysis suggests is needed to reach a 10^{-5} contrast level.

The simulations I present in this section make a simplifying assumption where we allocate the entire error budget for coronagraphy to come exclusively from piston errors on the OTE segments. In reality, there will be static surface errors on all segments, OTE misalignments like modeled here, science instrument-internal WFE contributions, WFS&C residuals and dynamic WFE that all contribute to the overall WFE budget. However, we can show that even when taking such internal WFEs into account, the total WFE requirement for $c_t = 10^{-5}$ is still less stringent than the required WFE from published JWST optical performance models. In this case, subtracting the WFE internal to the NIRCcam A module, which is 54 nm rms as modeled in WebbPSF, from the total requirement of 187 nm rms, we are left with a requirement of roughly 130 nm rms. Thus we conclude that the OTE alignment will not be limiting the contrast.

An additional point of interest is the differential sensitivity of the individual segments. While previous work has shown that placing segments with a higher static rms surface polishing error behind the secondary support struts optimizes the coronagraphic performance (Perrin et al. 2012), as one would intuitively assume, tolerancing with the PASTIS model allows us to quantify that. Since the individual JWST segments show a factor of four in the range they span in terms of their tolerable WFE for a particular target contrast, this shows which segment control needs to be emphasized by the alignment team. This quick analysis confirms that JWST will support the contrast values of the coronagraphs, but if there is ever a drive to optimize the performance further, for example better alignments in combination with aggressive post-processing, the tolerancing tools I developed could be used for an in-depth examination of the OTE requirements.

Starting from the original implementation of analytical images in order to calculate the pair-aberrated contrast influence (Lebouilleux et al. 2018b), the simulations in this section were used to establish the semi-analytical PASTIS matrix calculation (Laginja et al. 2019). This formal extension allows us to use existing E2E simulators for the purpose of creating a contrast influence matrix, the PASTIS matrix, and perform a statistical tolerancing analysis for a segmented, coronagraphic instrument. In addition, the work on the NIRCcam coronagraph on JWST proved that the tolerancing concept with PASTIS is portable to HCI instruments other than the 36-segment telescope concept used by Lebouilleux et al. (2018b), and that the sensitivity results are commensurable with the budgeted WFE on the NIRCcam A module, for a contrast of 10^{-5} or higher. The work in this section was an important milestone toward more realistic simulations of yet another segmented coronagraph, the HiCAT testbed, presented in the following section. The ultimate goal here is to prepare a full experimental segment-level tolerancing analysis, as presented in Chap. 4.

3.3 Contrast sensitivities to segment aberrations on the HiCAT testbed

Adapted from:

Predicting contrast sensitivity to segmented aperture misalignment modes for the HiCAT testbed

I. Laginja, R. Soummer, L. M. Mugnier, L. Pueyo, J.-F. Sauvage, L. Lebouilleux, L. Coyle, J. S. Knight, M. D. Perrin, S. D. Will, J. Noss, K. J. Brooks, J. Fowler

Proc. SPIE 11443, 114433J (2020)

With its intermediate-level contrast performance compared to envisioned missions like LUVOIR, and with a segmented mirror permanently in the optical beam, the HiCAT testbed provides excellent testing grounds for coronagraph sensitivity studies. Such activities are supported by the excellent software infrastructure surrounding the testbed, which contains a full optical simulator for optical modeling of the

bench, and a testbed emulator that enables numerical simulation of the hardware behavior. In practice, these are tightly linked: the top-level control software contains an interface that can either establish a link to the hardware testbed located in the Makidon laboratory, or fire up the testbed emulator linked to the optical simulator, which is distributed over a private GitHub repository. One of the main advantages of such an architecture is that experiments can be prepared without needing testbed access at least to a point at which there are no syntax errors in the software anymore.

Exploiting this setup, we use the HiCAT testbed emulator to implement experiments on the testbed's optical simulator. It allows us to simultaneously debug the control software for the experiments, as well as to carry out a performance assessment of the HiCAT coronagraph with respect to segment misalignments on the segmented DM included in the simulations. The intended goal is to define the co-phasing requirements of the segmented DM for any given contrast level on HiCAT, and to prepare the testbed infrastructure for the experimental validations presented in the following chapter. We first give a brief overview of the HiCAT testbed, present the testbed emulator, then describe the optical layout of the coronagraph on the bench including the active components, and finally we show the results of the simulations conducted with the HiCAT optical simulator.

3.3.1 The HiCAT project and experiment emulation setup

The HiCAT testbed (High-contrast imager for Complex Aperture Telescopes; N'Diaye et al. 2013, 2014, 2015a; Lebouilleux et al. 2017a, 2016; Moriarty et al. 2018; Soummer et al. 2018) is a demonstrator for coronagraphy with segmented apertures and integrated WFS&C systems (see Chap. 1.2.3). Aiming to investigate and improve technologies for LUVVOIR-type space telescopes of the future, it works in a mid-contrast (10^{-6} – 10^{-8}) regime in ambient air. The workhorse coronagraph configuration for HiCAT is an Apodized Pupil Lyot Coronagraph (APLC; N'Diaye et al. 2015b, 2016; Zimmerman et al. 2016). Since the apodizers on HiCAT are mounted on easily interchangeable bonding cells, it is easy to perform fast changes between different designs, or mount a high-quality flat mirror in the apodizer pupil plane to work as a classical Lyot coronagraph (CLC). For active WFS&C, HiCAT contains two 952-actuator Boston Micromachines DMs (Cornelissen et al. 2010), one located in-pupil, and one out-of-pupil, together with an IrisAO segmented DM (Helmbrecht et al. 2013). HiCAT is therefore truly segmented with the ability to add real co-phasing wavefront errors and introduce temporal drifts for dynamical studies.

HiCAT testbed emulator and controls

The high-level testbed control system of HiCAT is unique due to its dual-mode operation setup. Coded purely in Python, the same code base is used to control either the actual testbed hardware or to use the optical simulator for HiCAT, emulating the physical behavior of the testbed. The latter includes a full simulation of the hardware control interfaces, which makes it a complete “synthetic testbed”. In practice, this means that the *exact* scripts that are used to run testbed simulations can also be run on the actual testbed, allowing us to first test experiments on the simulator, and then work right out of the box on the hardware. This works with identical commands to the motor controllers, DMs, and all hardware components, using the same pipeline for data processing and producing the same output data products. It is this “emulated testbed” that we use to produce the results presented in this section.

The overall control system and software architecture is object-oriented and modular, and is hosted on GitHub, while also deploying automatic testing and continuous integration. The hardware controls have been abstracted in the public CATKit Python package (Noss et al. 2021a), which provides the interface to all our hardware components, for example from Boston Micromachines, IrisAO, Thorlabs, Newport and others. The HiCAT optical simulator uses the POPPY Python Fourier optics toolkit (Perrin et al. 2016a, 2012), mixing Fraunhofer and Fresnel models, including a fast semi-analytic coronagraph propagation at high resolution. The simulator was originally constructed from the theoretical testbed optical design and later refined with experimental calibrations and measured optical alignments. It enables model-based control algorithms with a Jacobian calculated on the simulator, for example for pair-wise probing (Groff et al. 2016) and stroke minimization (Mazoyer et al. 2018a,b; Pueyo et al. 2009).




Figure 3.10: Overlapping pupils in the HiCAT segmented CLC configuration used for the PASTIS numerical experiments, produced by the optical simulator. The entrance pupil mask (light green) traces the outline of the IrisAO, preventing the illumination of areas outside of the controllable segments. The Lyot stop (yellow circle) is sized such that its edges stay within the controllable outline of the IrisAO (light green). The entrance pupil diameter D_{pup} is defined as the circumscribed circle around the IrisAO (dashed white).

HiCAT optical configuration for numerical validation of PASTIS

While the HiCAT APLC is designed to provide a superior performance on a segmented aperture compared to the simpler CLC, the pupil plane apodization of this coronagraph causes a lot of the aperture segments to be highly concealed (see Soummer et al. 2018, Fig. 7). This has a direct impact on the segment-level tolerancing as described by the PASTIS model (Laginja et al. 2021, Sec. 6). Moreover, the pupil apodization and the FPM filtering display competing effects in the tolerancing. From a practical stand point, setting up HiCAT in a full segmented APLC operation mode requires to achieve an accurate alignment of the various pupil planes, especially the apodizer optic with respect to the segments of the IrisAO DM. To allow for an initial characterization of the coronagraph performance in the simpler CLC mode, the testbed will have no apodizer installed for a while, which is why we choose to perform and prepare the experimental validation of PASTIS on a CLC configuration of HiCAT. Having the IrisAO installed puts HiCAT into a “segmented CLC mode”, whose different pupils in simulation are plotted in Fig. 3.10. This includes a non-circular pupil mask that traces the segmented IrisAO outline (light green area), and its circumscribed circle is what we refer to as the pupil diameter D_{pup} (dashed white circle). HiCAT hosts one in-pupil and one out-of-pupil 1k Boston continuous DM, an FPM with a radius of $8.52 \lambda/D_{pup}$ and a circular, unobscured Lyot stop. Its diameter D_{LS} (dashed black line) is equivalent to the inscribed circle of the IrisAO (projected in the Lyot plane), keeping its edges within the controllable area of the segmented mirror (and the hexagonally outlined pupil mask), which puts its size at 81% of D_{pup} . With the WFS&C strategy on HiCAT, using




Figure 3.11: HiCAT optical configuration used for PASTIS experiments, here shown with transmissive optics for simplicity. The entrance pupil is a custom-shaped mask tracing the outline of the segmented DM in a consecutive pupil plane. Of the two continuous deformable mirrors, DM1 is in a pupil plane, and DM2 is located out-of-pupil. The focal-plane mask (FPM) and Lyot stop (LS) build the classical Lyot coronagraph setup.

pair-wise estimation and stroke minimization, the previously used unsegmented version of this CLC setup (with a circular pupil instead of the outline mask, with diameter D_{pup}) achieved a contrast of 4×10^{-8} on hardware, in monochromatic light at 640 nm, and 2×10^{-7} in 10% broadband light, in a 360° DH from $6-11 \lambda/D_{LS}$ (where the outer working angle is defined by the highest spatial frequency controllable by the continuous DMs).

This optical testbed layout, used for the simulated experiments, can be seen in Fig. 3.11. The hexagonally outlined pupil mask lies in a conjugate pupil plane with the segmented DM, continuous DM1 and the LS. The continuous DM2 is mounted, and in the current case simulated, 30 cm out of pupil in order to correct for amplitude aberrations with the Talbot effect. Both the FPM and the science detector are located in focal planes. The segments of the IrisAO are each controllable in piston, tip and tilt, with a maximum stroke of $5 \mu m$. Using a 14-bit controller, we assume close to perfect linear actuators (Helmbrecht & Juneau 2007), resulting in a smallest control step per actuator of 0.3 nm, excluding any noise. This control step will drive the contrast level we tolerance the segmented WFE limits to in Sec. 3.3.2.

With this segmented CLC setup on the HiCAT emulator, without any WFS&C and with a flat segmented DM, the simulated average contrast in the DH area from $6-11 \lambda/D_{LS}$ is 1×10^{-5} . In order to place the coronagraph floor of the testbed into a higher contrast regime, we deploy an iterative WFS&C loop that uses pair-wise sensing to estimate the E-field in each iteration, followed by a control step with both DMs as calculated by the stroke minimization algorithm. The IrisAO is kept at its best flat position throughout. We choose to stop this control loop after 10 iterations, which is when we reach an average DH contrast of 5.7×10^{-8} , comparable to the best contrast HiCAT can reach on the real testbed in an unsegmented CLC configuration³, and in monochromatic light. The DM solutions and focal plane image from this simulated experiment are displayed in Fig. 3.12. The WFS&C solution shown in Fig. 3.12 is included in the static coronagraph contribution of the PASTIS model (Laginja et al. 2021, Eq. 12), which allows us to work around an improved best contrast solution compared to the coronagraph without deploying a DH algorithm. This sets our coronagraph floor that we use in the emulated PASTIS experiments on HiCAT in Sec. 3.3.2 to an initial $c_0 = 5.7 \times 10^{-8}$.

3.3.2 Results on the HiCAT emulator

Simulating sensitivity experiments on HiCAT with the PASTIS tolerancing model allows us to estimate the segmented DM WFE requirements for decided target contrast levels, and to prepare the experiment software. We are using the “synthetic testbed” mode of the hardware controls to write the experiment code, and the results from the optical simulator will allow us to draw conclusions about the sensitivity of the HiCAT CLC to IrisAO misalignments. These numerical experiments include all preparation steps needed for the anticipated hardware runs and they are expected to run “out of the box” on the testbed

³The segmented CLC configuration later reached a best contrast of 2.5×10^{-8} on hardware, see also Chap. 4.




Figure 3.12: DM surface commands applied to continuous DM1 (in-pupil, left) and DM2 (out-of-pupil, center) to obtain the 360° DH solution from $6\text{--}11 \lambda/D_{LS}$ on the HiCAT emulator (right), calculated by 10 iterations of pair-wise sensing and stroke minimization, which yields an average contrast of 5.7×10^{-8} . The segmented DM in this setup is kept at its best flat position throughout. We include these DM solutions in the definition of our coronagraph when working with PASTIS.

hardware. In this section, we present the results of these emulated validation experiments. We use the testbed configuration described in Sec. 3.3.1, using a monochromatic light source at 640 nm. Before running an experiment, we apply the DM solutions for the DH shown in Fig. 3.12 in order to use the DH contrast from that stroke minimization solution as the coronagraph floor, initially $c_0 = 5.7 \times 10^{-8}$.

In the following, we present the results of three numerical experiments for the validations of the PASTIS tolerancing model: (1) measuring a simulated “experimental” PASTIS matrix and validating the PASTIS forward model (Eq. 2.9) with a “hockey stick curve” experiment (see Luginja et al. 2021, Fig. 4), (2) measuring the cumulative contrast of the modes obtained from the emulated PASTIS matrix by tolerancing all modes to a uniform contrast contribution, and (3) performing a simulated Monte Carlo experiment for the validation of the calculated segment requirements, where we propagate random WFE maps drawn from the tolerancing prescription for independent segments calculated with PASTIS to measure the resulting DH average contrasts.

The first simulated experiment is the measurement of a PASTIS matrix, which will be used to calculate the mode and segment requirements for a given target contrast. As described in Sec. 3.3.1, the least significant bit (LSB) of the IrisAO controller allows for a minimal movement of 0.3 nm of a single IrisAO segment, barring any noise. In order to minimize limitations by the LSB, we choose here a conservative target contrast of $c_t = 10^{-6}$, which results in a standard deviation for the segment requirements larger than 4 nm. Since these tolerances are drawn from a zero-mean distribution, some random WFE realizations will still be truncated to zero due to the LSB limit, especially when taking into account additional controller noise, but with larger standard deviations we are increasing the fraction of realizations above that limit, and the effect becomes negligible.

PASTIS matrix measurement and validation

The PASTIS matrix is a pair-aberrated influence matrix, linking segment aberrations to the average contrast in the coronagraphic DH. Having the scalar quantity of the average DH contrast as its objective, it has the advantage that no prior knowledge of the electric field is required in order to construct it. This means that measuring an experimental PASTIS matrix is much simpler than measuring an empirical electric field Jacobian, since its calculation does not introduce overheads or estimation errors that usually come with E-field estimation methods (e.g., measuring probe images in the pair-wise estimator). One aspect to consider is that the measurement time for the PASTIS matrix scales roughly with the square of the number of segments (n_{seg}) divided by two, rather than linearly with the number of segments:

$$n_{meas} = \frac{n_{seg}^2 + n_{seg}}{2}, \quad (3.1)$$

where n_{meas} is the total number of measurements required for the construction of the PASTIS matrix. We divide by 2 because the matrix is symmetrical, and include the matrix diagonal by adding n_{seg} . On the




Figure 3.13: *Left:* PASTIS matrix for HiCAT as measured in an emulated experiment. Each entry represents the differential contrast contribution of each aberrated segment pair, normalized to the aberration amplitude. The matrix is symmetric, and its diagonal shows the impact on contrast by the individual segments, which is used in the independent segment tolerancing. *Right:* Geometry of the IrisAO segmented DM on HiCAT and the segment numbering used in this paper. The 37 segments are numbered starting at 0 for the center segment, to 36 in the outer ring.

37-segment HiCAT pupil, this requires only $n_{meas} = 703$ measured images.

In the presented numerical experiment, we obtain the PASTIS matrix with the HiCAT emulator. We constrain ourselves to a local piston mode with an amplitude of 10 nm rms WFE for the calibration aberration of the PASTIS matrix. Other modes are possible, for example tip/tilt, or a combination of local segment aberrations, but they are not considered in this paper. Following the semi-analytical approach, we first calculate the contrast matrix by aberrating pairs of segments and record the resulting DH average contrast. We then use Eqs. 2.16 and 2.17 to calculate the PASTIS matrix shown in Fig. 3.13, left. The PASTIS matrix is symmetric, with its diagonal describing the impact on contrast by the individual segments, which is used in the independent segment tolerancing. There are some blue, negative streaks in the matrix with a very low change of contrast for particular segment pairs, which correspond to adjacent segments in the pupil.

To validate the PASTIS forward model using the PASTIS matrix in Eq. 2.9, we generate random and independent segment phase aberrations over the entire segmented DM, represented by the vector \mathbf{a} , scale them to a global rms WFE and propagate them with the PASTIS matrix. In parallel, we also apply this aberration to the HiCAT segmented DM in simulation and measure the resulting DH contrast. Since one particular rms WFE over the total pupil can be realized with many different individual segment configurations, we average over the contrast values from 10 different realizations at each global rms WFE value. The result of this simulation is shown in Fig. 3.14. We observe that the two propagators show very good accordance, and more so in the small aberration regime up to ~ 2 nm rms, just above the coronagraph floor, even if the error between them grows only marginally beyond that. We can clearly see the curve flatten out toward the left, where they are limited by the coronagraph floor, and the two solutions start diverging from each other at a global WFE value of around 100 nm rms, which lies way beyond the small-WFE approximation regime on all segments.

We proceed with an eigendecomposition of the PASTIS matrix (see Chap. 2.3) and calculate its eigenmodes, shown as the PASTIS modes in Fig. 3.15. The modes are ordered from highest to lowest eigenvalue, indicating their comparative impact on the DH average contrast in their natural normalization. In a similar fashion to the eigenmodes in Chap. 2.3, the respective contrast impact is reflected in the spatial frequency content of each mode. The highest-influence eigenmodes in the first two rows of Fig. 3.15 all show checkerboard-like patterns in the unconcealed areas of the pupil, while the low-impact modes approximate low-order Zernikes, which the coronagraph is inherently less sensitive to. In the following section, we scale all modes to a uniform contrast contribution between them.




Figure 3.14: Validation of the PASTIS matrix by propagating the same segmented WFE maps both with the semi-analytical PASTIS matrix from Fig. 3.13 (solid blue) and with the emulated HiCAT testbed (dashed orange). The curve flattens out to the left, at the coronagraph floor c_0 , and shows linear behavior at larger WFE, giving it its hockey stick-like shape. The two propagators show very good accordance in the small aberration regime right above the contrast floor, between ~ 0.6 and 4 nm rms over the entire segmented mirror.

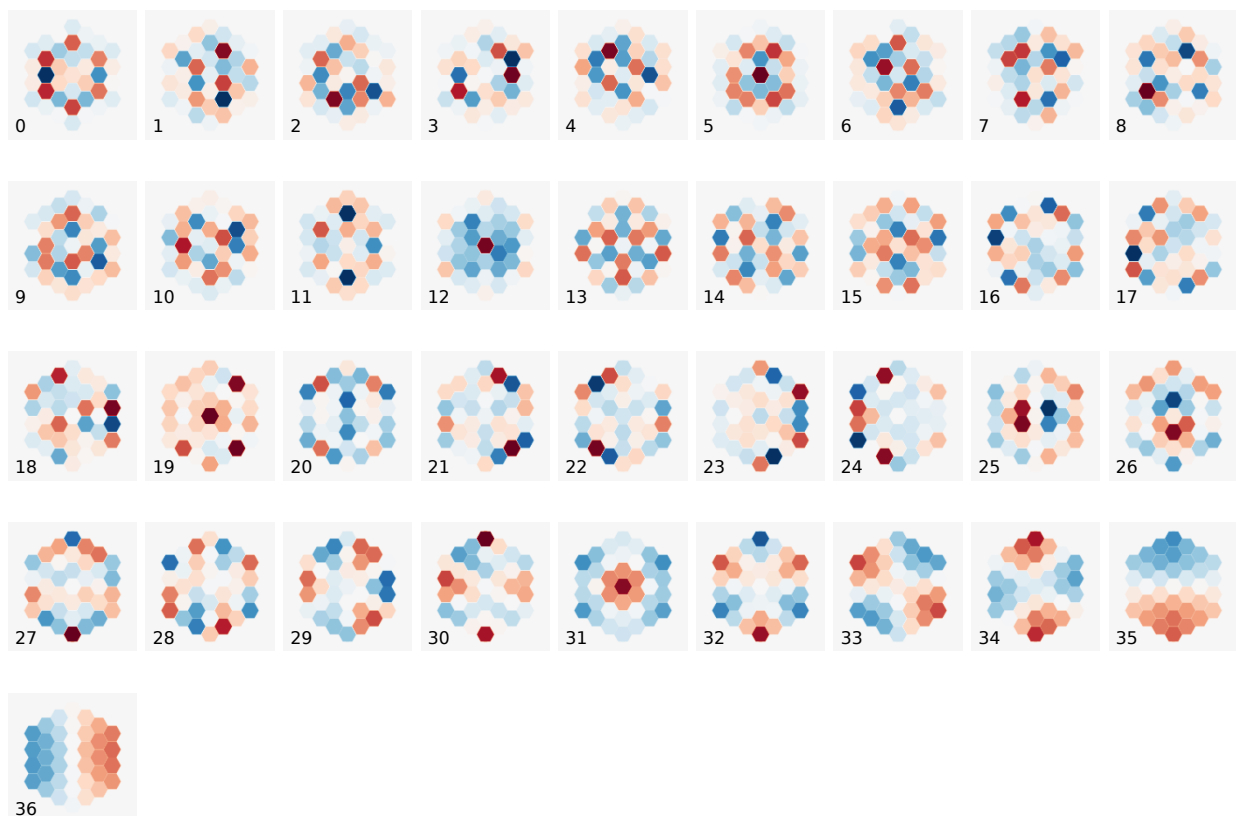


Figure 3.15: All simulated PASTIS modes for HiCAT with a classical Lyot coronagraph, for local piston aberrations, sorted from highest to lowest eigenvalue. The modes are unitless, showcasing the relative scaling of the segments to each other, and between all modes. They gain physical meaning when multiplied by a mode aberration amplitude b_p in units of wavefront error or phase. Their relative impact on final contrast is given by their eigenvalues.




Figure 3.16: *Left:* Mode requirements as calculated with Eq. 2.26 for a uniform contrast contribution per mode to a target contrast of 10^{-6} . *Right:* Cumulative contrast plot for the uniform mode requirements shown left, calculated both with the PASTIS model equation (solid blue) as well as measured with the HiCAT emulator (dashed orange). The emulator result shows a better accordance with the PASTIS model at lower mode index. Note how neither line starts at the coronagraph floor because the mode with index 0 already adds a contrast contribution on top of the baseline contrast.

Validation of mode tolerances

The PASTIS modes in Fig. 3.15 form an orthonormal mode basis - each of them contributes to the overall contrast additively, without influence from the other modes. This can be used to define error budgets based purely on these optical modes (see Chap. 2.3.2). In the present example, we choose that the PASTIS modes should contribute uniformly to the total contrast, in which case we can calculate the mode tolerances for a particular target contrast with Eq. 2.26. The resulting mode requirements for a target contrast of $c_t = 10^{-6}$ are shown in Fig. 3.16, left. To validate the assumption of modes that are additive in contrast, we run a simulated experiment to measure the cumulative contrast of the tolerated PASTIS modes. For this, we multiply the modes by their respective requirement, apply them cumulatively to the IrisAO and measure the resulting DH average contrast at each step (Fig. 3.16, right). The cumulative measurements with the HiCAT emulator follow the general expected linear shape, although some mode contributions seem to overshoot its predicted contrast contribution slightly. These over-contributing modes then seem to be compensated by weighted modes that do not influence the contrast quite as much as intended, displaying a periodic error pattern, and reaching a final contrast 5% (5×10^{-8}) above the target contrast of 10^{-6} .

Validation of independent segment tolerances

To fully validate the PASTIS tolerancing model, we calculate segment-level requirements and probe them with the HiCAT simulator. In cases where the segments can be assumed to be independent from each other, as is the case for an IrisAO, we can calculate individual segment requirements with Eq. 2.36 as a function of the target contrast. While the overall level of WFE requirements will be highly influenced by the Fourier filtering of the FPM, the different segments display a differential tolerance between them, see Fig. 3.17, left. These individual segment requirements will be highly influenced by pupil features of the optical system. Looking at their spatial distribution in the HiCAT pupil, we can see in Fig. 3.17, right, that the segments of the outer ring have more relaxed requirements than the two inner rings and the center segment. This is caused, in large part, by the Lyot stop, which is covering a large fraction of the segments in the outer ring because it is undersizing the pupil, which can be seen in Fig. 3.10. The segment-level WFE requirements displayed in Fig. 3.17 present a statistical description of the allowable segment-level WFE if a target contrast of 10^{-6} is to be maintained as a statistical mean over many misalignment states of the segmented DM. As long as the change of the segment-level WFEs on the DM follow zero-mean normal




Figure 3.17: Independent segment requirements for HiCAT simulations as calculated with Eq. 2.36 for a target contrast of 10^{-6} as a line plot (left) and as a spatial map (right). These numbers are the standard deviation of the tolerable WFE rms on each segment if the target contrast is to be met. The requirement range spans from 4 to 6.5 nm, with a clear jump in the outermost ring, which is highly concealed by the Lyot stop (see Fig. 3.10). The same segment numbering is used as in Fig. 3.13, right.

distributions whose standard deviations per segment are described by the numbers in Fig. 3.17, the target contrast will be recovered as the statistical mean over many such realizations.

In order to confirm this assumption, we proceed by running a simulated Monte Carlo experiment, producing 1000 different WFE aberration patterns on the segmented DM and recording the propagated average DH contrast. Taking data for 1000 realizations is doable on the hardware in a time frame of about one hour, so we expect to retrieve a histogram with about the same accuracy when performing this experiment on the testbed in the near future. The tolerances in Fig. 3.17 are the prescription as to how to draw these random WFE realizations: each segment-level WFE on segment k , in a single random WFE map (vector \mathbf{a}), is drawn from its own zero-mean normal distribution with a standard deviation of μ_k from Eq. 2.36. This means that one random HiCAT WFE map is composed of 37 independent normal distributions with a mean of zero, and a standard deviation of μ_k , which then gets applied to the IrisAO on the HiCAT emulator and propagated through the coronagraph to measure the DH contrast. The distribution of measured average contrast values is shown in Fig. 3.18. The resulting figure corresponds to a Gaussian distribution with a mean of 1.07×10^{-6} and a standard deviation of 2.57×10^{-7} , marked in the plot with dark red lines. To interpret the results in Fig. 3.18, we remember that PASTIS provides analytical expressions to derive the expected mean contrast (Eq. 2.31) and variance (Eq. 2.32) from a distribution calculated with a set of segment requirements. Apart from the PASTIS matrix M , what is needed to calculate these quantities is the segment covariance matrix C_a , which in the case of independent segments as presented here is a simple diagonal matrix made of the segment requirement variances, μ_k^2 , which we take from Fig. 3.17. Then, Eq. 2.32 yields an analytical standard deviation of 2.35×10^{-7} and the mean is, as expected, the target contrast value 10^{-6} . These are marked with yellow lines in Fig. 3.18. We observe that both the measured mean as well as the variance are slightly higher than the analytically calculated values from the PASTIS matrix. This discrepancy could be attributed to a drifting coronagraph floor from simulated jitter, noise and WFE realizations in the emulator. These sources of WFE vary from one exposure to the next and would bias the coronagraph floor that is captured in each individual DH measurement. This renders the baseline contrast c_0 different from what was used for the calculation of the segment tolerances. Since this effect would also influence the measured PASTIS matrix, this error could bias the contrast measurements in the MC experiments towards higher values. We identify this as a crucial point to address for the experimental validations on hardware, where the baseline contrast is known to fluctuate over longer periods of time.

Overall, our numerical experiments on the HiCAT emulator present a successful implementation of simulations of the PASTIS model for a specific high-contrast instrument, the HiCAT testbed. We measured a simulated PASTIS matrix and validated it by comparing its propagation results with measurements from the synthetic testbed. We decomposed the matrix into independent optical modes that we toleranced uniformly and cumulatively to a target contrast of 10^{-6} . Finally, we calculated segment-level WFE re-




Figure 3.18: Monte Carlo validation experiment on the HiCAT emulator to validate the independent segment error budget shown in Fig. 3.17, for a target contrast $c_t = 10^{-6}$. Each segmented WFE map draws from 37 zero-mean distributions with an individual standard deviation per segment, μ_k . The expected mean contrast and standard deviation of this distribution, as calculated by Eqs. 2.31 and 2.32, are $\langle c \rangle = 10^{-6}$ (the target contrast) and 2.35×10^{-7} (dashed and dotted yellow lines). The simulated distribution has a mean of 1.07×10^{-6} and a standard deviation of 2.57×10^{-7} , both of which are slightly larger than the predicted values. This is likely to stem from a combination of a drifting coronagraph floor due to simulated image jitter. A mitigation for this effect is presented in Chap. 4.

quirements under the assumption of independent segments and validated them with a numerical Monte Carlo experiment, measuring the contrast from randomly drawn segmented WFE maps as prescribed by the derived requirements.

3.4 Conclusions

The PASTIS tolerancing analysis is useful to determine the coronagraphic sensitivity on a particular telescope to aberrations. In order to put such results into perspective, comparative studies between different instruments provide insight into how such sensitivities vary with different telescope parameters like coronagraph type, number of segments, telescope diameter, etc. To enable such studies, I created a tool that can do the full PASTIS analysis from calculating a pair-aberrated matrix to computing the respective WFE requirements per aberration mode or segment. This tool is a Python package called “PASTIS” that provides a unified framework for this procedure.

In order to perform such an analysis on a particular telescope or instrument, the package contains clean interfaces to connect various numerical simulators for the purpose of generating PASTIS matrices. In addition to this flexible interface, the package itself implements its own optical simulators for a number of telescopes, including LUVOIR A with three APLC designs, LUVOIR B with a VC, and the monolithic HabEx with a VC. They contain the capability to include various aberration sources that can be controlled like a DM, with their influence functions being either global Zernikes, ripples, local Zernikes, or custom local aberrations, as well as actuator-based DMs. The package uses OOP for the implementation of these simulators as well as the calculation of the PASTIS matrices, and new releases are regularly published with Zenodo and distributed over GitHub. The development of this package was a major contributor to later work presented in Chap. 5.

The modular interface allows us to perform the tolerancing analysis not only on the package-internal

telescope simulators, but also on coronagraphs based in the open-source Python package WebbPSF, that is JWST and RST. The work that went into this enabled the development of the semi-analytic PASTIS matrix calculation, which makes the matrix more physically accurate than the previously used fully analytical computation. Moreover, performing the analysis on a CLC of NIRCcam opened up the interface to allow for any E2E simulator to be used, for example the HiCAT simulator.

The work presented in this chapter was necessary to prepare the experimental validations of the PASTIS tolerancing on hardware, presented in Chap. 4. Connecting the PASTIS analysis package with the HiCAT simulator lead to successfully performed simulated experiments. From these results, we deduct that the segmented mirror needs to be stable to $\sim 4\text{--}6$ nm in order to enable a mean contrast of 10^{-6} . Our goal is to verify these numbers on hardware. Considering the synthetic testbed nature of the HiCAT emulator, the experiments were bound to directly work on hardware. Indeed they did, with only the fine-tuning of the experimental parameters left before producing the results in the following chapter. This shows the importance of numerical simulators: they can provide fairly accurate proof-of concept results that are a necessary step before investing into hardware demonstrations. More than just optical simulations, I want to stress the advantages of full numerical testbed emulation. Preparing experiments on a laboratory setup requires lots of work both on the hardware itself, as well as significant time dedicated to developing the experiment control codes and their debugging. When this is done with direct hardware access, the testbed is blocked both for potential parallel work on the hardware, as well as any sort of data acquisition. By having a numerical emulator that is installable on any machine, the implementation and debugging process can be offloaded away from the primary lab setup up to the point where there are no bugs left in the code, and preliminary results are achieved numerically. This ensures an optimization in the usage time of the testbed, as direct access is only needed when actual measurements are being made, or to test software integrations. In particular for the case of HiCAT, this allows for several projects being integrated and tested at the same time on separate instances of the testbed emulator, while the time on the laboratory setup itself is optimally exploited.

While running the simulated experiments on the HiCAT testbed emulator, we anticipated certain challenges before moving to the real hardware. The limitations imposed by the IrisAO will define how well we can control the aberration modes we introduce. While this segmented DM is known to have close to perfect linear behavior, the modes applied to it will stem from open-loop calibrations, which might include some errors. Furthermore, the least significant bit introduced by the IrisAO controller will prevent us from aberrating a segment with an amplitude smaller than ~ 1 nm when taking noise into account. This will lead to drawing from imperfect normal distributions when creating random WFE maps with the prescription of the segment requirements, truncating any small aberrations to zero, which might skew the results of the Monte Carlo analysis. However, we expect these errors to be small if the target contrast is chosen high enough compared to the coronagraph floor, which is a direct motivation to improve the overall performance of HiCAT, to which I dedicated a major part of my PhD. Another challenge will be the stability of the DH solution that we adopt into our coronagraph, setting the contrast floor we perform our experiments at, as the DH contrast is directly influenced by fluctuating environmental parameters like temperature and humidity. While one solution is to compensate for these fluctuations by upgrading the WFS&C performance to run faster, another one is to account for these fluctuations in the analytical model of PASTIS, which is addressed in Chap. 4.

4

WAVEFRONT STABILITY TOLERANCING IN EXPERIMENTAL VALIDATIONS

With the results of the HiCAT simulations from the previous chapter being completed in early December 2020, I obtained the first experimental results of the tolerancing validations already a few weeks later, just before Christmas 2020. While the initial experimental data revealed some control issues with the IrisAO segmented deformable mirror on HiCAT, those were solved right after the holidays and I was analyzing my initial testbed results only a couple of weeks later, in early 2021. With some additional work on understanding the connection between different experimental parameters and their interpretation, I published the results in a peer-reviewed journal paper, which builds this chapter (Laginja, I. et al. 2022).

The fast turnaround time in going from simulations to hardware experiments is enabled by the integrated hardware emulator described in the previous chapter, which controls the actual hardware components when run on the testbed computer, and otherwise uses a suite of simulated testbed components. This allowed me to use an installation of the full HiCAT testbed infrastructure on my local machine to prepare my experiments from Europe, while a remote desktop connection into the Makidon lab enabled me to perform the experiments myself from across the globe. The credit that needs to be given to this operational setup cannot be emphasized enough. While there was already an interest in these tools for 24/7 operations during full on-site work and before I moved to France for the second half of my PhD, the Coronavirus pandemic underlined the advantage of being able to control the testbed fully remotely, and let it run autonomously for several hours without supervision.

The experimental validations of the PASTIS tolerancing model on HiCAT gave us an opportunity to study the coronagraphic contrast sensitivity to segmented DM misalignments in particular. Since such perturbations are only one of many contributors to WFE variations over time which in turn influence the contrast, it was necessary to investigate to what extent different sources of WFE can be separated in the analysis. The result is a slight modification in the analytical model that now relates everything to a drifting contrast difference rather than absolute contrast. A major goal in this was to compare analytical tolerancing results to observations in the lab, from which I deduced a segmented DM error budget for various contrast levels on the HiCAT testbed.

Wavefront tolerances of space-based segmented telescopes at very high contrast: Experimental validation

I. Leginja, J.-F. Sauvage, L. M. Mugnier, L. A. Pueyo, M. D. Perrin, J. Noss, S. D. Will, K. J. Brooks, E. Por, P. Petrone, R. Soummer

Published in
A&A 658, A84 (2022)

4

ABSTRACT

Context. The detection and characterization of Earth-like exoplanets (exoEarths) from space requires exquisite wavefront stability at contrast levels of 10^{-10} . On segmented telescopes in particular, aberrations induced by cophasing errors lead to a light leakage through the coronagraph, deteriorating the imaging performance. These need to be limited in order to facilitate the direct imaging of exoEarths.

Aims. We perform a laboratory validation of an analytical tolerancing model that allows us to determine wavefront error requirements in the $10^{-6} - 10^{-8}$ contrast regime, for a segmented pupil with a classical Lyot coronagraph. We intend to compare the results to simulations, and we aim to establish an error budget for the segmented mirror on the High-contrast imager for Complex Aperture Telescopes (HiCAT) testbed.

Methods. We use the Pair-based Analytical model for Segmented Telescope Imaging from Space (PASTIS) to measure a contrast influence matrix of a real high-contrast instrument, and use an analytical model inversion to calculate per-segment wavefront error tolerances. We validate these tolerances on the HiCAT testbed by measuring the contrast response of segmented mirror states that follow these requirements.

Results. The experimentally measured optical influence matrix is successfully measured on the HiCAT testbed, and we derive individual segment tolerances from it that correctly yield the targeted contrast levels. Further, the analytical expressions that predict a contrast mean and variance from a given segment covariance matrix are confirmed experimentally.

4.1 Introduction

The search for Earth-like exoplanets (exoEarths) and potential signs of life in the form of atmospheric biomarkers is a very exciting field of today's astronomy. However, it requires a tremendous improvement in imaging capabilities compared to what we can currently achieve, in order to capture the few photons coming from a planet buried in the blinding light of its nearby host star. Instruments will need to reach contrast levels (planet to star flux ratios) of at least 10^{-10} , at a separation of only ~ 0.1 arcsec, or less, from the star (The LUVOIR Team 2019). Telescopes providing these capabilities will require large collecting areas, and they will most likely be realized with segmented primary mirrors, both in space and on the ground. Currently, the favored method to achieve these extreme high-contrast levels are dedicated instruments called coronagraphs that strongly attenuate the on-axis star light while preserving the off-axis planet light as much as possible (Guyon et al. 2006). These instruments are very sensitive to residual wavefront aberrations, which generate speckles of light in the imaging focal plane that can be mistaken for planets. This is why coronagraphy needs to be combined with wavefront sensing and active control (WFS&C) (Mazoyer et al. 2018a,b; Groff et al. 2016) to create a zone of deep contrast in the final image, a dark hole (DH). These ambitious goals will be achieved from space by missions such as the Habitable Exoplanet Observatory (HabEx; Gaudi et al. 2019) and the Large UV Optical InfraRed Surveyor (LUVOIR; The LUVOIR Team 2019; Bolcar 2019) which are currently under consideration by the NASA Astro2020 Decadal Survey, with the Nancy Grace Roman Space Telescope (RST; Krist et al. 2015) working toward shorter-term demonstrations at more moderate contrast levels ($10^{-7} - 10^{-9}$) with a monolithic primary mirror.

The extreme contrast levels that are needed for the detection of exoEarths require excellent stability against wavefront errors (WFEs) over a range of temporal and spatial frequencies (Pueyo et al. 2019; Coyle et al. 2019b; Feinberg et al. 2017). While some of these can be actively controlled with a WFS&C system, or do not have a large impact on the contrast due to robust coronagraph designs, aberration modes to which the instrument is very sensitive must be held to a minimal level (Nemati et al. 2020; Juanola-Parramon et al. 2019b; Nemati et al. 2017b). Typically, it is enough to control these misalignment modes on the timescales of the WFS&C system. In this paper, we focus on the segment-related aberrations due to segment misalignments, which are the main contributors to WFE in the mid-spatial-frequency regime (Douglas et al. 2019; Juanola-Parramon et al. 2019b; Moore & Redding 2018). Various technology solutions are being developed toward this goal (Coyle et al. 2020; Hallibert et al. 2019; Coyle et al. 2018; East et al. 2018; Stahl 2017; Stahl et al. 2015). Multiple hardware efforts are underway to provide laboratory demonstrations of the systems anticipated to be installed on future large observatories. While other testbeds are tackling the problem on monolithic apertures (Potier et al. 2020a; Patterson et al. 2019; Sidick et al. 2015), the two that have been focusing on segmented apertures are the High-contrast imager for Complex Aperture Telescopes (HiCAT) testbed at the Space Telescope Science Institute (Soummer et al. 2018) and the High Contrast Spectroscopy Testbed for Segmented Telescopes (HCST) at Caltech (Llop-Sayson et al. 2020).

The tolerancing problem of segmented high-contrast instruments has been previously addressed with the Pair-based Analytical model for Segmented Telescope Imaging from Space (PASTIS; Laginja et al. 2021, 2020, 2019; Lebouilleux et al. 2018b,a, 2017b). It is an analytical model that calculates the average contrast in the DH caused by pupil-plane segment misalignments, using a simple matrix multiplication. Central to this model is the PASTIS matrix M , which describes the contrast contributions of an aberrated segment pair. When we combine it with a covariance matrix, C_a , that describes the thermo-mechanical behavior of the segments, the PASTIS model can be used to calculate the expected mean contrast and its variance for a particular instrument, over many segmented aberration states, with analytical equations. This allows us to fully describe the statistical response of a segmented coronagraph to segment-level cophasing errors. Additionally, an eigendecomposition of either matrix allows us to write the contrast as a sum of separate contributions and then to invert the problem: instead of calculating the expected mean contrast given some aberrations, we can now statistically determine tolerances that lead to a particular mean contrast target. This leads to a quantitative tool for instrument design that provides the means to calculate statistical limits on the segmented aberration modes.

In this paper, we use the semi-analytical development of the PASTIS model (Laginja et al. 2021, 2019) to

measure an experimental PASTIS matrix on the HiCAT testbed (Soummer et al. 2018; Moriarty et al. 2018; Leboulleux et al. 2016, 2017a; N'Diaye et al. 2015a, 2014, 2013) by replacing the images usually provided by an end-to-end simulator with laboratory measurements. We then use this experimental PASTIS matrix for further analysis, meaning the validation of the instantaneous forward model for the calculation of the average DH contrast and WFE tolerancing analysis, and compare the results to a simulated case. First, we aim to provide a general experimental validation of the PASTIS model, showing that it is feasible to measure a PASTIS matrix on hardware, and that we can calculate correct segment-level tolerances in the mid-contrast regime ($10^{-6} - 10^{-8}$), compatible with the HiCAT performance as of today. This is demonstrated in Sec. 4.4 at 2.5×10^{-8} , which is the contrast floor considered in this paper. Our tolerancing process is demonstrated for this value of contrast, situated at an intermediate point between the expectations of the James Webb Space Telescope (JWST) and LUVOIR. In Sec. 4.4, it translates into tolerancing values of a few nanometers, compatible with the hardware limitations of the bench (resolution of deformable mirrors and fast fluctuations in the optical system). We eventually demonstrate the sensitivity to a contrast change of around a few 10^{-7} . Second, we compare the results obtained with testbed measurements to results obtained with a purely simulated PASTIS matrix, and we analyze the sensitivity to model errors. And third, we establish an error budget for the segmented deformable mirror (DM) on HiCAT, for various contrast levels, and present quantitative results on the required WFE stability of the segmented mirror.

In Sec. 4.2 we recall the most important points about the PASTIS forward model and its inversion, from which we derive the results for WFE tolerancing. We also include a simple extension for the treatment of a drifting coronagraph floor that is not attributed to the segmented mirror. In Sec. 4.3 we describe the HiCAT project and the testbed configuration used for the presented experiments. In Sec. 4.4 we show the tolerancing results and their validations performed on the HiCAT testbed, and we compare them to a simulated case of the PASTIS matrix on HiCAT. Finally, in Secs. 4.5 and 4.6 we discuss our results and report our conclusions.

It should be noted that the main figure of merit used by the PASTIS model is the spatially averaged intensity in the DH, normalized to the peak of the direct image, which is what we refer to as “contrast” throughout this paper; it depends on the particular state of the instrument, and notably that of the segments. We also stress that we differentiate between this spatially averaged DH intensity, the “average DH contrast”, and a statistical mean (expected value) of this averaged contrast over many optical propagations, the statistical “mean contrast”.

4.2 PASTIS tolerancing model and extension to a drifting coronagraph floor

While segmented aberrations have a direct impact on the focal plane response, they are not the only source of time fluctuations in the contrast. For a laboratory validation, the environmental conditions (humidity, temperature, vibrations) evolve with time and contribute to opto-mechanical deformations of the testbed, eventually translating into slowly evolving optical aberrations and contrast drift. We need to take this contrast drift, which is not due to the segmented mirror, into account to be able to isolate the effects coming from the segmented mirror alone.

In the following section, we present a brief summary of the PASTIS tolerancing model (Laginja et al. 2021), and expand it to include a drifting coronagraph floor arising from time-dependent aberrations from sources other than the segmented mirror.

We model the phase ϕ in the pupil plane of a segmented high-contrast instrument as:

$$\phi(\mathbf{r}, t) = \phi_{DH}(\mathbf{r}) + \phi_{ab}(\mathbf{r}, t) + \phi_s(\mathbf{r}), \quad (4.1)$$

where ϕ_{DH} is a static best-contrast phase solution, usually produced by a DH algorithm and applied to a pair of DMs. The term ϕ_{ab} is the phase produced by time-dependent aberrations in the system, ϕ_s is the phase caused by segment-induced aberrations, \mathbf{r} is the pupil plane coordinate and t the time variable. Under the assumption of the small aberration regime for ϕ_s and ϕ_{ab} , and assuming that ϕ_{DH} is static, we discard any cross-terms between ϕ_s and ϕ_{ab} as they would create third and fourth order terms in contrast, while we limit ourselves to a second order model. We can then express the electric field in the pupil

with two terms: the first is a time-dependent term that includes the best-contrast phase solution with an additional aberrating phase drift, and the second is independent from time and contains the segmented perturbations:

$$E(\mathbf{r}, t) = P(\mathbf{r}) e^{i\phi(\mathbf{r}, t)} \simeq P'(\mathbf{r}, t) + iP(\mathbf{r}) e^{i\phi_{DH}(\mathbf{r})} \phi_s(\mathbf{r}), \quad (4.2)$$

where P is the pupil function, and $P'(\mathbf{r}, t) = P(\mathbf{r}) \exp[i(\phi_{DH}(\mathbf{r}) + \phi_{ab}(\mathbf{r}, t))]$. Applying a linear coronagraph operator, C , that represents the propagation of the electric field in the high-contrast system (i.e., Fourier transforms and mask multiplications) to the expression given in Eq. 4.2, we obtain the coronagraphic intensity distribution in the image plane with $|C\{E(\mathbf{r}, t)\}|^2$. The average contrast is then given by averaging over the DH area, indicated by $\langle \dots \rangle_{DH}$.

It was previously shown that the average DH contrast can always be expressed as a quadratic function of a segmented phase perturbation under an appropriate change of variable (Laginja et al. 2021, Eqs. 4 and 11), which eliminates the linear cross-term, and leaves us with separate square transformations of the two terms in Eq. 4.2. Concretely, in this paper we model the contrast floor with a contribution from the static DM phase, and aberrations introduced by environmental changes of the testbed, which cause a drift in the contrast as a function of time t ,

$$c_0(t) = \langle |C\{P'(\mathbf{r}, t)\}|^2 \rangle_{DH} = \langle |C\{P(\mathbf{r}) e^{i(\phi_{DH}(\mathbf{r}) + \phi_{ab}(\mathbf{r}, t))}\}|^2 \rangle_{DH}. \quad (4.3)$$

Representing optical aberrations on a segmented telescope with local (per-segment) Zernike modes, we can expand the phase aberrations on the segmented pupil $\phi_s(\mathbf{r})$ in the second term of Eq. 4.2 as a sum of segment-level polynomials (Laginja et al. 2021, Eq. 1), with its decomposition on such a basis denoted as \mathbf{a} . Following the development of the original PASTIS model, we can express the average contrast in the coronagraphic DH as a matrix multiplication (Laginja et al. 2021, Eq. 9), which makes our assumed model:

$$c(t) = c_0(t) + \mathbf{a}^T M \mathbf{a}, \quad (4.4)$$

where $c(t)$ is the spatial average contrast in the DH, $c_0(t)$ the coronagraph floor (i.e., the average contrast in the DH in the presence of the best-contrast phase $\phi_{DH}(\mathbf{r})$ and of the variable phase aberrations $\phi_{ab}(\mathbf{r}, t)$, but in the absence of segment misalignments ϕ_s), M is the symmetric PASTIS matrix of dimensions $n_{seg} \times n_{seg}$ with elements m_{ij} , \mathbf{a} is the aberration vector of the local Zernike coefficients on all discrete n_{seg} segments and \mathbf{a}^T its transpose. We can see that the contrast floor $c_0(t)$ is dominated by the DH phase solution $\phi_{DH}(\mathbf{r})$, with an additional variation introduced by the time-dependent $\phi_{ab}(\mathbf{r}, t)$. The matrix M itself contains a constant term added by $e^{i\phi_{DH}(\mathbf{r})}$, and the segmented aberrations induced by $\phi_s(\mathbf{r})$. Following the pairwise aberrated approach explained previously (Laginja et al. 2021, Eq. 10), the matrix elements can thus generally be expressed as:

$$m_{ij} = \langle C\{P(\mathbf{r}) e^{i\phi_{DH}(\mathbf{r})} Z(\mathbf{r} - \mathbf{r}_i)\} C\{P(\mathbf{r}) e^{i\phi_{DH}(\mathbf{r})} Z(\mathbf{r} - \mathbf{r}_j)\}^* \rangle_{DH}. \quad (4.5)$$

By defining the differential contrast as our objective quantity that is independent of time t :

$$\Delta c = c(t) - c_0(t), \quad (4.6)$$

we render the right-hand side of Eq. 4.4 time-independent, which allows us to isolate the effects imposed by segment cophasing errors, defined by the vector \mathbf{a} .

Each PASTIS matrix element m_{ij} represents the contrast contribution to the DH average contrast c_{ij} by each aberrated segment pair in the pupil, formed by segments i and j . Once the matrix is established, we can calculate its eigenmodes \mathbf{u}_p and eigenvalues λ_p by means of an eigendecomposition. The total number of optical (PASTIS) modes, n_{modes} , is equal to the total number of segments, n_{seg} . Since the eigenmodes are orthonormal and diagonalize M , the DH contrast can be written as the sum of separate contributions of each mode, and each eigenvalue is the contrast sensitivity of the corresponding mode (Laginja et al. 2021, Eq. 22):

$$\Delta c = \sum_p^{n_{modes}} b_p^2 \lambda_p, \quad (4.7)$$

where b_p is the amplitude of the p^{th} mode.

With the PASTIS matrix, M , representing the optical properties of the segmented coronagraph, and Eq. 4.4 giving the instantaneous average DH contrast, c , for a given aberration vector \mathbf{a} , M can be combined with any given segment covariance matrix, C_a , to calculate the statistical mean and variance of the average DH contrast (Laginja et al. 2021, Eqs. 31 and 32):

$$\langle \Delta c \rangle = \text{tr}(MC_a), \quad (4.8)$$

where tr denotes a trace and

$$\text{Var}(\Delta c) = 2 \text{tr}[(MC_a)^2]. \quad (4.9)$$

While Eq. 4.8 does not make any assumptions about the statistics of the vector \mathbf{a} , Eq. 4.9 is true when \mathbf{a} follows a Gaussian distribution, which is an assumption used throughout this paper. The two equations above allow us to calculate these two integral quantities directly from the optical properties of the instrument, described by M , and the mechanical correlations of the segments, captured by C_a , which can be obtained from thermo-mechanical modeling of the observatory.

For all C_a , Eq. 4.8 can be expressed as

$$\langle \Delta c \rangle = \sum_{p=1}^{n_{\text{modes}}} \sigma_{b_p}^2 \lambda_p, \quad (4.10)$$

where σ_{b_p} are the standard deviations of the optical mode amplitudes, in the diagonalized basis of the PASTIS matrix M . This leads to an inversion of the problem where we set a differential target contrast Δc , for which we want to derive WFE tolerancing limits in terms of standard deviations for the segments, or modes. For the special case of a diagonal C_a , which means that the individual segments are statistically independent, a similar expression to Eq. 4.10 can be deduced from Eq. 4.8 for the standard deviations of the segment amplitudes, σ_{a_k} :

$$\langle \Delta c \rangle = \sum_{k=1}^{n_{\text{seg}}} \sigma_{a_k}^2 m_{kk}, \quad (4.11)$$

where the m_{kk} are the diagonal elements of the PASTIS matrix. This equation can be used to specify the standard deviation for each segment: denoting $\mu_k = \sigma_{a_k}$, we can choose that every segment contributes equally to the contrast, which yields a specification of segment amplitude standard deviations of (Laginja et al. 2021, Eq. 36):

$$\mu_k^2 = \frac{\langle \Delta c \rangle}{n_{\text{seg}} m_{kk}}. \quad (4.12)$$

While Eqs. 4.8 and 4.9 allow us to analytically calculate the expected mean contrast of a segmented coronagraph and its variability, for mechanical properties described by C_a , Eq. 4.12 provides a way to determine individual segment tolerances for a particular target differential contrast Δc that is to be maintained over a set of observations.

4.3 The HiCAT project and experimental setup

The HiCAT testbed (Soummer et al. 2018; Moriarty et al. 2018; Leboulleux et al. 2016, 2017a; N'Diaye et al. 2015a, 2014, 2013) is dedicated to a LUVOIR-type coronagraphic demonstration with on-axis segmented apertures¹. The project is targeting experiments in ambient conditions that can happen before demonstrations in a vacuum, for example at the Decadal Survey Testbed (DST; Patterson et al. 2019) located at the Jet Propulsion Laboratory. The ultimate performance goal of such testbeds is to demonstrate a contrast of 10^{-10} in the lab. While this goal can only be achieved in an environmentally stable vacuum chamber, the HiCAT testbed is aiming for 10^{-8} , limited by its coronagraph performance and environmental conditions. The work on HiCAT intends to provide a system-level analysis of a high-contrast instrument that includes various sensors and controllers. Ultimately, the planned coronagraph for HiCAT operations is an Apodized

¹https://exoplanets.nasa.gov/internal_resources/1186/




Figure 4.1: HiCAT optical configuration used for PASTIS experiments, here shown with transmissive optics for simplicity. The entrance pupil is a custom-shaped mask that traces the outline of the segmented DM in a consecutive pupil plane (see also Fig. 4.2). Of the two continuous DMs, DM1 is in a pupil plane, and DM2 is located out-of-pupil in order to control both phase and amplitude. The FPM and LS form the CLC setup. The pupil planes (PP) and focal planes (FP) are marked.

Pupil Lyot Coronagraph (APLC; N’Diaye et al. 2016; Zimmerman et al. 2016; N’Diaye et al. 2015b) that includes apodizers manufactured using carbon nanotubes. Since the various apodizer designs are mounted on easily interchangeable bonding cells, a high-quality flat mirror can be swapped in to use a classical Lyot coronagraph (CLC). We use the CLC setup (Sec. 4.3.1) for the experiments in this paper, supported by an active WFS&C loop to improve the DH contrast beyond the initial static solution caused purely by the coronagraphic masks (Sec. 4.3.2). An IrisAO (Helmbrecht et al. 2013) segmented DM is utilized as the segmented telescope simulator on the testbed to introduce segment-level WFEs for the tolerancing validation with PASTIS.

The high-level testbed control system of HiCAT is coded in Python and uses the Numpy (Oliphant 2006; Van der Boekel et al. 2011), Matplotlib (Hunter 2007; Caswell et al. 2020), Astropy (The Astropy Collaboration et al. 2013; The Astropy Collaboration et al. 2018; The Astropy Collaboration 2018), SciPy (Virtanen et al. 2020), scikit-image (van der Walt et al. 2014), pandas (McKinney 2010; Reback et al. 2020), imageio (Silvester et al. 2020), photutils (Bradley et al. 2020), HCIPy (Por et al. 2018), Poppy (Perrin et al. 2012) and CatKit (Noss et al. 2021b) packages. The software infrastructure emulates all testbed processes including the control of all hardware components, which means that all experiments can be prepared offline, without any access to the actual hardware (Laginja et al. 2020).

4.3.1 Classical Lyot coronagraph as static setup

While the HiCAT APLC is designed to provide a superior performance on a segmented aperture compared to the simpler CLC, the pupil plane apodization of this coronagraph causes a lot of the aperture segments to be highly concealed (Soummer et al. 2018, Fig. 7). HiCAT has been operated as a segmented CLC since the fall of 2020, and we used this testbed configuration to perform the experimental validations of PASTIS, which allows us to image the entire segmented aperture (unobstructed 37 segments).

The defining optical elements of HiCAT with a CLC are a noncircular pupil mask, an IrisAO PTT111L 37-element hexagonally segmented DM (Helmbrecht et al. 2016, 2013), two Boston Micromachines 952-actuator microelectro-mechanical (MEMS) “kilo-DMs” (Cornelissen et al. 2010), a focal-plane mask (FPM), a Lyot stop (LS) and a science detector. A schematic of the optical testbed layout used for the presented experiments can be seen in Fig. 4.1. The pupil mask traces the hexagonally segmented IrisAO outline, and we use its circumscribed diameter as the pupil diameter, D_{pup} . This is slightly undersized with respect to the diameter of the IrisAO itself to limit the beam to the controllable area of the segmented DM. The first Boston continuous DM is located in a pupil plane, and the second one is out-of-pupil, at a distance of 30 cm from the pupil plane DM, in order to control amplitude in the WFS&C process. The FPM has a diameter of $8.52 \lambda/D_{pup}$, with λ (no subscript) the central wavelength of the bandpass. The LS is a circular, unobscured mask with a diameter D_{LS} of 79% of the size of D_{pup} , as projected in the Lyot plane. An overlay of all




Figure 4.2: *Left:* Overlapping pupils in the HiCAT segmented CLC configuration, and their diameters, used for the PASTIS experiments, projected in the pupil plane of DM1 (simulated image). The entrance pupil mask (bright green, undersized polygon shape) traces the outline of the IrisAO (dark green polygon shape) with a 97% undersizing factor, preventing the illumination of areas outside of the controllable segments. The entrance pupil diameter D_{pup} (white dashed line) is defined as the circumscribed circle around the undersized pupil mask. The LS (yellow ellipse) is sized such that its edges stay within the controllable outline of the IrisAO. Since HiCAT uses reflective optics, the resulting beam foreshortening along the x axis results in all pupils being optically “squished” along the x direction. This is true for all optics and is immediately visible in this figure as the yellow LS surface is slightly elliptical with respect to the circle denoting its nominal diameter D_{LS} (dashed black line). *Right:* Measured pupil image on hardware with a detector located in a pupil plane before the two continuous DMs and the LS, showing the IrisAO segments and the pupil mask outlining the segmented DM. We note the slightly undersized outline, which results in somewhat irregular hexagons at the edges, especially noticeable on the six corner segments.

relevant pupils can be seen in Fig. 4.2.

Previously, the IrisAO had been used with the CLC for experiments on coronagraphic focal plane wavefront sensing (WFS) on a segmented aperture (Lebouilleux et al. 2020), but with different mask sizes and a fully circular entrance pupil. The installation of the IrisAO on the current CLC setup was performed in late 2020. The segments of the IrisAO segmented DM are each controllable in piston, tip and tilt, with a maximum stroke of $5 \mu\text{m}$ on each of the three actuators mounted on the back side of each segment. The segmented DM initially saw an open-loop flatmap calibration (Helmbrecht et al. 2016) with a 4D Fizeau interferometer, which yielded a calibrated surface error of 9 nm root-mean-square (rms) (Soummer et al. 2018, Fig. 3). This was improved upon after installing the IrisAO on the testbed, where a finer, closed-loop flat-map calibration was performed with differential optical transfer function (dOTF) phase retrieval (Codona 2012; Codona & Doble 2012).

The average contrast for the currently used CLC setup, in an annular DH from $6\text{--}10 \lambda/D_{LS}$ with flattened DMs is $\sim 1 \times 10^{-5}$ in monochromatic light at 638 nm . In order to place the coronagraph floor of the testbed into a higher contrast regime, we deploy an iterative WFS&C loop described in the following section.

4.3.2 Active wavefront sensing and control for an improved DH

The WFS&C strategy used on HiCAT to improve the monochromatic DH contrast in an annular DH deploys an iterative approach of pair-wise probing (Groff et al. 2016; Give’on et al. 2011) to estimate the electric field, and stroke minimization (Mazoyer et al. 2018a,b; Pueyo et al. 2009) for control. The outer

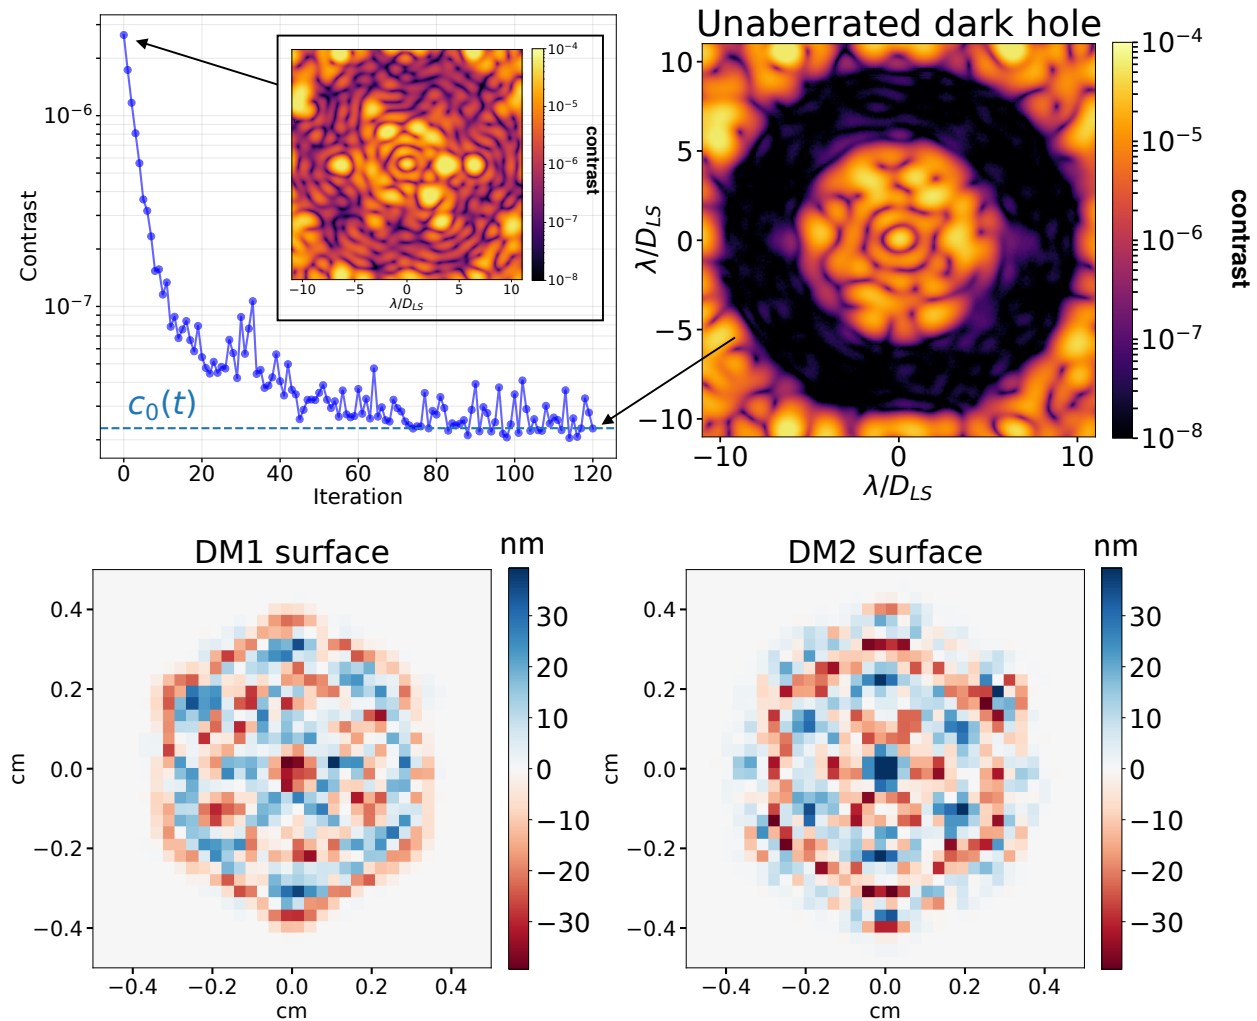


Figure 4.3: *Top:* 120 iterations of pair-wise sensing and stroke minimization (left), starting from a pre-modulated coronagraphic image shown in the embedded plot. The loop converges after 70–80 iterations and yields an average contrast of 2.5×10^{-8} from 6–10 λ/D_{LS} , using a monochromatic source at 638 nm (right). The two DH images are shown on the same scale. *Bottom:* DM surface commands applied to continuous DM1 (in-pupil, left) and DM2 (out-of-pupil, right) for the best-contrast DH solution in the top right at iteration 120. The segmented DM in this setup is statically set to its best flat position throughout the control loop.

working angle of 10 λ/D_{LS} is defined by the highest spatial frequency controllable by the two continuous DMs, and the IrisAO DM is kept at its best flat position throughout. In order to avoid a local minimum, we first dig a larger DH at moderate contrast before launching the loop on a 6–10 λ/D_{LS} DH as seen in Fig. 4.3. After 70–80 iterations, the contrast performance converges to 2.5×10^{-8} , with variations on the order of 2×10^{-8} during the WFS&C sequence. The continuous DM commands that create the final DH, as well as the convergence plot and the final DH image, are shown in Fig. 4.3. The DM surface commands shown in Fig. 4.3 are applied at the beginning of each experiment presented in Sec. 4.4, making them a part of the static coronagraph contribution $c_0(t)$ as described by Eq. 4.3. This setup sets our nominal coronagraph floor that we use in the PASTIS experiments on HiCAT in Sec. 4.4 to an initial 2.5×10^{-8} , and it is drifting without active control during the experiments. The aberrations we target with the PASTIS tolerancing model in this paper are the segmented WFEs, \mathbf{a} , introduced by the IrisAO DM on top of this static best-contrast solution, and independent of the contrast drift.

4.4 Experimental validation of segmented tolerancing on the HiCAT testbed

In the following section, we present the results of several experiments for the hardware validation of the PASTIS tolerancing model. We measure an experimental PASTIS matrix and compare it to simulations, and we confirm the instantaneous PASTIS forward model in Eq. 4.4. We measure the deterministic optical mode contrast given by Eq. 4.7 and then validate the statistical segment tolerances calculated from the experimental PASTIS matrix with Eq. 4.12 by performing Monte Carlo experiments and comparing them to results from Eqs. 4.8 and 4.9. We use the testbed configuration described in Sec. 4.3, and we describe how to correct our measurements for the drifting contrast floor $c_0(t)$ in Sec. 4.4.1. Before running an experiment, we apply the continuous DM solutions for the DH shown in Fig. 4.3, putting the testbed initially at $c_0(t_0) = 2.5 \times 10^{-8}$. In Sec. 4.4.1, we emphasize the need to refine our forward model to account for the slow contrast drift on the testbed, as introduced in Sec. 4.2. At this level of performance, this drift is the main limitation and has to be corrected numerically. In Sec. 4.4.3, we compute the tolerancing in terms of segment allocations corresponding to delta contrast values of a few 10^{-7} , as limited by the uncorrected fast fluctuations we see in our data in Fig. 4.5.

4.4.1 PASTIS matrix measurement and deterministic forward model validation

Measurement method

The PASTIS matrix is a pair-wise influence matrix, linking segment aberrations to the differential average contrast in the coronagraphic DH. The total number of intensity measurements needed for the construction of an experimental PASTIS matrix is:

$$n_{meas} = \frac{n_{seg}(n_{seg} + 1)}{2}. \quad (4.13)$$

Indeed, the matrix being symmetrical, we measure only the non-repeating permutations of segment pairs including the matrix diagonal. On the 37-segment HiCAT pupil this requires $n_{meas} = 703$ measurements.

The relation between each pair-wise aberrated contrast measurement and the PASTIS matrix elements is given by (Laginja et al. 2021, Eq. 15):

$$c_{ij}(t) = c_0(t) + a_i^2 m_{ii} + a_j^2 m_{jj} + 2a_i a_j m_{ij}, \quad (4.14)$$

where a_i is the WFE amplitude on segment i . We used the same calibration aberration amplitude $a_c = a_i = a_j$ that is put on each individual segment in the measurement of the contrast matrix (c_{ij}). The calibration of the M matrix is thus obtained by the measurement of the contrast from pushing a pair of segments (i, j). Since the natural testbed contrast is evolving with time, this measurement must be corrected for the coronagraph floor persisting at that same time, $c_0(t)$, which can be easily remeasured for each $c_{ij}(t)$. The expression for the diagonal matrix elements then becomes:

$$m_{ii} = \frac{c_{ii}(t) - c_0(t)}{a_c^2}. \quad (4.15)$$

This makes the PASTIS matrix diagonal entirely independent of time and the coronagraph floor, and it describes the contrast contribution of each individual segment to the DH. Ideally, $c_{ii}(t)$ and $c_0(t)$ are measured at the same time t ; in reality, they are measured within a short time of each other, which needs to be faster than the occurring drifts and can thus be assumed to be simultaneous. Since this corrects the diagonal matrix elements for the coronagraph floor at the time of their measurement, we now want to make the off-diagonal elements depend on the already calibrated diagonal PASTIS matrix elements m_{ii} , rather than the uncalibrated diagonal measurements of the contrast c_{ii} . We can easily solve Eq. 4.14 for the off-diagonal PASTIS matrix elements:

$$m_{ij} = \frac{c_{ij}(t) - c_0(t)}{2a_c^2} - \frac{m_{ii} + m_{jj}}{2}. \quad (4.16)$$

In this way, each matrix element m_{ij} gets calibrated with an appropriate, time-dependent measurement of $c_0(t)$, which makes the entire PASTIS matrix time-independent.




Figure 4.4: Examples of pair-wise aberrated DH images. *Top:* Simulated WFE maps of aberrated segment pairs on the IrisAO segmented DM. The LS cuts off most of the segments in the outer ring of the segmented pupil. The segment numbering is as indicated in Fig. 4.6, left. *Bottom:* DH images from the testbed with resulting fringes from pair-wise aberrated segments during the PASTIS matrix measurement shown in the top row. The DH extent is from $6\text{-}10 \lambda/D_{LS}$, indicated with the dashed circles, and the images are displayed in the same range as the unaberrated DH in Fig. 4.3. The pair (3|4) is made of adjacent segments, which produces low-spatial-frequency fringes in the DH. This leads to an overall decrease in the contrast contribution, as is confirmed by their respective entry in the PASTIS matrix in Fig. 4.6, with a blue entry right next to the matrix diagonal.

Matrix measurement

We use this to measure an experimental PASTIS matrix on the HiCAT testbed. We constrain ourselves to a local piston mode with a WFE amplitude of $a_c = 40$ nm rms for the calibration aberration of the PASTIS matrix. Other modes are possible, for example tip and tilt or a combination of local segment aberrations, but they are not considered in this paper. An aberration of 40 nm rms on a single segment of a 37-segment pupil translates to a global aberration of 6.6 nm rms, while two such aberrated segments cause a global WFE of 9.3 nm rms. We can see in Fig. 4.7 that this puts a_c significantly off the knee around the coronagraph floor, which was discussed as an optimal regime for a_c in Laginja et al. (2021). This was done in order to increase the signal-to-noise ratio (SNR) in the fringe images during the matrix calibration, while simultaneously not increasing a_c too much along the linear aberration regime indicated in Fig. 4.7.

We sequentially aberrate pairs of segments by applying the calibration aberration a_c to each segment to measure $c_{ij}(t)$, and then flatten the IrisAO DM and record the coronagraph floor $c_0(t)$ for the same iteration. Examples of pair-wise aberrated DHs causing fringe patterns are displayed in Fig. 4.4, and the evolution of the unaberrated coronagraph floor during the PASTIS matrix acquisition is shown in Fig. 4.5.

We then use Eqs. 4.15 and 4.16 to calculate the elements m_{ij} and construct the experimental PASTIS matrix M^{exp} shown in Fig. 4.6, middle. The PASTIS matrix is symmetric, with its diagonal describing the impact on the contrast by the individual segments. There are some negative streaks in the matrix, colored blue in the figure. Such negative matrix elements $m_{ij} < 0$ are interference terms that reduce the contrast loss, meaning that the sum of intensities when pushing segments i and j individually gives a worse contrast than pushing them at the same time. This phenomenon is strongly correlated to the spatial frequency of the fringe pattern created by the segment pair. For a high spatial frequency (distant




Figure 4.5: Contrast $c_0(t)$ during the PASTIS matrix acquisition. After each pair-wise aberrated DH measurement, we flatten the IrisAO segmented DM to measure the drift in the coronagraph floor over the course of the experiment. We subtract this off our data in order to perform an analysis on the differential contrast Δc . This open-loop contrast degrades gradually over time, note the linear scale. The measured contrast values range from 2.5×10^{-8} to 8×10^{-8} during the course of the experiment, but the difference between adjacent measurements is initially on the order of 2×10^{-8} , rising to 1×10^{-7} later on, which is sufficient for our proposed calibration method. The total duration of the experiment is 45 minutes.

segments), the contrast degradation is averaged over the DH and is therefore minimized. For a low spatial frequency (close segments), the contrast can be degraded or improved due to the spatial configuration of the DH with respect to the fringe pattern. We also show a simulated PASTIS matrix M^{sim} in Fig. 4.6, right, calculated with the same calibration aberration per segment of $a_c = 40 \text{ nm}$ WFE rms, but without any WFE or measurement noise in the optical system. We can see that the general morphology of the simulated and experimental matrices is the same – in particular, it is the same segment pairs that show the highest and lowest contrast contribution in the image plane, relatively speaking. The experimental matrix is noisier though, and it has a slightly higher overall amplitude. Here, we want to show that it is feasible to directly measure an experimental PASTIS matrix, which will represent the real optical system more accurately, and compare this to results obtained with the simulated matrix.

Contrast model validation

To validate the instantaneous PASTIS forward model, we compare the contrast for a segmented phase error calculated with Eq. 4.4, against the contrast measured for the same segmented phase map applied to the IrisAO on the testbed. For a range of rms WFE values, we generate a random segmented phase map \mathbf{a} , and we scale it to a given global rms WFE. Then we evaluate Eq. 4.4 both with the simulated matrix M^{sim} and with the experimental matrix M^{exp} , and measure the resulting average DH contrast on the testbed. The results are plotted in Fig. 4.7. We observe that the results from the PASTIS equation using the experimental matrix (solid blue) show very good accordance with the testbed measurements (dashed orange), the curves overlap at a global rms WFE beyond 2 nm. The contrast calculated with the simulated matrix (solid green)




Figure 4.6: *Left:* Geometry of the IrisAO segmented DM on HiCAT and the segment numbering used in this paper, in the entrance pupil. The 37 segments are numbered, from 0 for the center segment to 36 in the outer ring. In the exit pupil, most of the outer ring segments are obscured by the LS (see Fig. 4.4). *Middle:* Experimental PASTIS matrix for HiCAT as measured on the testbed. Each entry represents the differential contrast contribution of each aberrated segment pair. The matrix is symmetric, and its diagonal shows the impact on the contrast by the individual segments. *Right:* Simulated PASTIS matrix for HiCAT, without any WFE or noise in the optical system. This matrix shows the idealized contrast contributions from each segment pair in a perfect system.

yields an equally accurate result compared to the hardware measurements. We can clearly see all curves flatten out toward the left, where they are limited by the coronagraph floor, producing a hockey stick-like shape.

4.4.2 Validation of mode contrast allocation

We proceed with an eigendecomposition of the experimental PASTIS matrix (Laginja et al. 2021, Sec. 3) and calculate its eigenmodes, shown as the optical PASTIS modes in Fig. 4.8. The modes are ordered from highest to lowest eigenvalue, indicating their comparative impact on the DH average contrast in their natural normalization. These segmented PASTIS optical modes for the HiCAT testbed with a CLC represent the modal contrast sensitivity of the instrument with respect to segment misalignments, with the sensitivity quantified by their respective eigenvalues. As shown previously, the lowest-impact modes (bottom of Fig. 4.8) are dominated by low-spatial-frequency components that are similar to discretized Zernike modes. The two modes with indices 35 and 36 in particular, the lowest-sensitivity modes, represent orthogonal tip and tilt modes over the entire pupil. High-impact modes (top of Fig. 4.8) display high-spatial-frequency content, mostly in the central area of the pupil which is unconcealed by the LS.

The PASTIS modes in Fig. 4.8 form an orthonormal mode basis, making them independent from each other - each of them contributes to the overall contrast without influence from the other modes, see Eq. 4.10. This can be used to define deterministic contrast allocations based purely on these optical modes (Laginja et al. 2021, Sec. 3.2). In the present example, we chose that each PASTIS mode should contribute uniformly to the total contrast, in which case we can calculate the exact mode weights for a particular target contrast:

$$\sigma_{b_p} = \sqrt{\frac{\langle \Delta c_t \rangle}{n_{modes} \cdot \lambda_p}}. \quad (4.17)$$

In accordance to the formalism laid out in Sec. 4.2, we make the mode tolerances independent of any given coronagraph floor by relating them to the differential target contrast Δc_t , displayed for a target contrast of $\Delta c_t = 10^{-6}$ in Fig. 4.9. To validate the assumption of a contrast that is a simple sum of separate mode contributions, we run an experiment to measure the cumulative contrast of the deterministically scaled PASTIS modes. For this, we multiply the modes by their respective uniform requirement, σ_{b_p} , apply them cumulatively to the IrisAO, measure the resulting DH average contrast at each step and subtract the simultaneously measured coronagraph floor $c_0(t)$ from the results (Fig. 4.10). We perform these propagations of the experimental eigenmodes both with the analytical PASTIS model in Eq. 4.4, using the experimentally measured matrix (solid blue), as well as with the HiCAT testbed (dashed orange). In Fig. 4.10, the




Figure 4.7: Validation of the deterministic forward model in Eq. (4.4) made by computing the contrast from the same segmented WFE maps with the experimental PASTIS matrix as in the middle panel of Fig. 4.6 (solid blue), with the simulated PASTIS matrix from the right panel of Fig. 4.6 (solid green), and with the HiCAT testbed (dashed orange). The curves flatten out to the left at the coronagraph floor, c_0 , and show linear behavior at increasing WFE, giving them their hockey-stick-like shape. The contrast calculation from the PASTIS equation with both matrices shows very good accordance with the testbed measurements; all three lines overlap at WFEs larger than 1 nm rms.

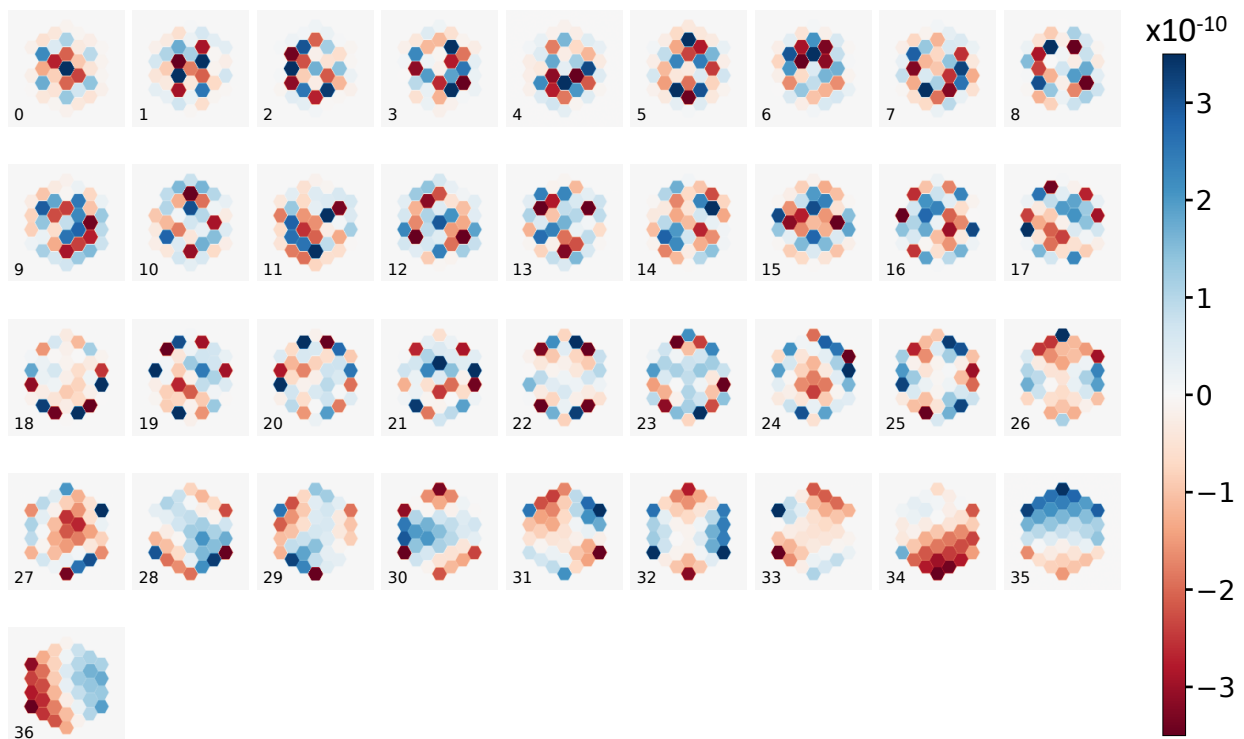


Figure 4.8: All experimental PASTIS modes for HiCAT with a CLC, for local piston aberrations, sorted from highest to lowest eigenvalue. The modes are unitless, showcasing the relative scaling of the segments to one another, and between all modes. They gain physical meaning when multiplied by a mode aberration amplitude, σ_{b_p} , in units of WFE or phase. Their relative impact on the final contrast is given by their eigenvalues.




Figure 4.9: Mode requirements as calculated with Eq. 4.17 for a uniform contrast contribution per mode to a differential target contrast of 10^{-6} . These are used as modal weights in the cumulative contrast measurement shown in Fig. 4.10.




Figure 4.10: Cumulative contrast plot for the uniform mode requirements shown in Fig. 4.9, calculated both with the PASTIS forward model in Eq. 4.4, with the experimental PASTIS matrix (solid blue), and measured on the HiCAT testbed (dashed orange). The analytical PASTIS calculation shows the perfect linear curve that the uniformly scaled modes would cause on an ideal testbed without noise or drifts. The testbed measurements show noisy behavior around the ideal curve, which is likely caused by segmented DM command uncertainties for small actuations and an insufficiently well-calibrated coronagraph floor.




Figure 4.11: *Left:* Independent segment requirements as calculated with Eq. 4.12 for target contrasts of $\Delta c_t = 5 \times 10^{-7}$ (top blue) and $\Delta c_t = 3 \times 10^{-7}$ (bottom orange). These numbers are the standard deviation of the tolerable WFE rms on each segment, if the target contrast is to be met as a statistical mean over many realizations of the segmented mirror. The range for the 5×10^{-7} target spans from 2.5 to 6 nm. Both curves show a clear jump in the outermost ring (starting with segment number 19), which is highly concealed by the LS (see Fig. 4.2). *Right:* Tolerance map for a target contrast of $\Delta c_t = 5 \times 10^{-7}$, shown in blue on the left. The global left-right asymmetries are due to a slightly offset LS, which does not impact the contrast performance at this level, but is visible in the segment sensitivities. The map for $\Delta c = 3 \times 10^{-7}$ is a scaled version of the one shown here.

cumulative measurements with HiCAT follow the general expected linear shape as displayed with the analytical PASTIS forward propagation using the experimental PASTIS matrix, but with some variations. These likely come from calibration errors in the segmented DM actuator influences for small displacements and contrast floor subtraction that is not accurate enough considering there will be no averaging effect during the comparatively short duration of this experiment (~ 5 min).

4.4.3 Statistical validation of independent segment tolerances

Experimentally calibrated segment tolerances

To fully validate the PASTIS tolerancing model for contrast stability, we calculate statistical segment-level requirements, from the experimental PASTIS matrix, for two target contrast values and measure their statistical contrast response with HiCAT. In this context, we use “tolerances” and “requirements” synonymously. In cases where the segments can be assumed to be independent from each other, as is the case for an IrisAO, we can calculate individual segment requirements (Laginja et al. 2021, Sec. 4.2) with Eq. 4.12 as a function of the differential target contrast. While the overall level of WFE requirements will be highly influenced by the Fourier filtering of the FPM, the different segments do not show uniform tolerance levels, as shown in Fig. 4.11, left. These individual segment requirements are highly influenced by pupil features of the optical system. Looking at their spatial distribution in the HiCAT pupil, we can see in Fig. 4.11 (right) that the segments of the outer ring have more relaxed requirements than the two inner rings and the center segment. This is caused, in large part, by the LS, which is covering a large fraction of the segments in the outer ring because it is undersizing the pupil, which can be seen in Fig. 4.2, left. The two sets of segment-level WFE requirements displayed in Fig. 4.11, left, represent a statistical description of the allowable WFE per segment if a delta target contrast of 5×10^{-7} or 3×10^{-7} is to be maintained as a statistical mean over many states of the segmented DM. We can observe how these two target contrast values yield segment requirements between 2.5 and 6 nm of WFE, which we know the IrisAO can reliably do; anything lower might lead to issues with the minimal stroke of the segmented DM and needs to be characterized in the future. Further, these particular delta contrast values are commensurable with the current testbed performance of HiCAT, staying just above the largest contrast fluctuations observed in

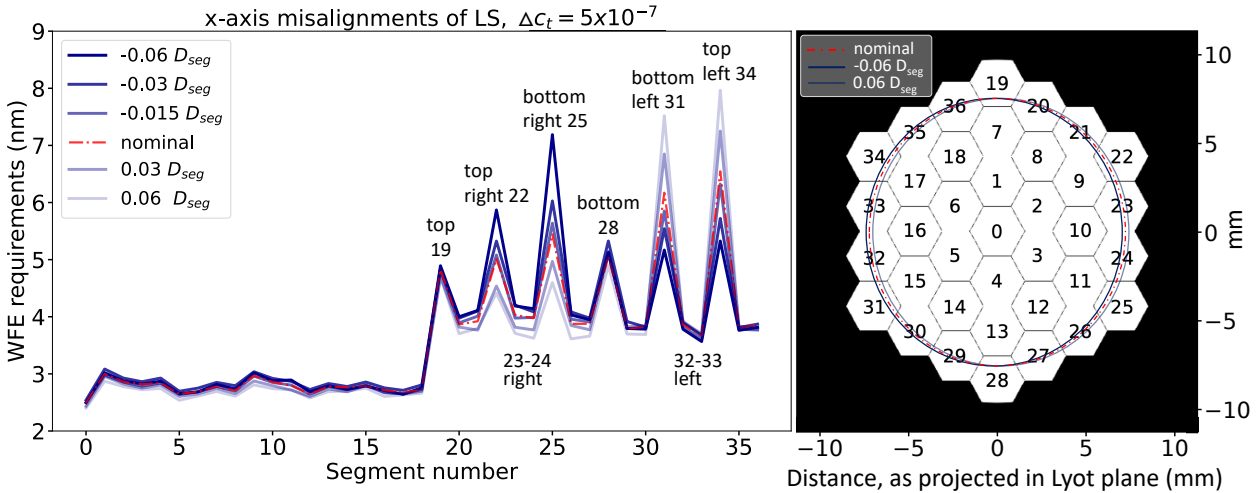


Figure 4.12: *Left:* Individual segment requirements for a target contrast of $\Delta c_t = 5 \times 10^{-7}$, measured for a range of x-axis misalignments of the LS. The legend is sorted in the same order as the lines for segment 25. The translation values are quoted in relation to the circumscribed segment diameter $D_{seg} = 1.4$ mm, projected into the Lyot plane. The labels indicate select segments in the outer ring. *Right:* Numbered segmented DM with overlaid LS positions. The nominal alignment is centered on the pupil (dashed-dotted red ellipse), with the dark and bright solid ellipses showing the maximum misalignments measured to either side. For details, see main text.

Fig. 4.5 ($\sim 1 \times 10^{-7}$), which are fast fluctuations that our model does not take into account.

As long as the components of the segment-level WFE on the DM follow independent zero-mean normal distributions whose standard deviations are given by the numbers in Fig. 4.11, the target contrast will be recovered as the statistical mean over many such realizations. The tolerance map in Fig. 4.11, right, shows a spatial representation of the segment-level standard deviations for $\Delta c = 5 \times 10^{-7}$; the map for $\Delta c = 3 \times 10^{-7}$ is a scaled version of the one plotted here and is not shown in this paper. The geometrical setup of the HiCAT pupil suggests a symmetry in the segment sensitivities along two axes, meaning for example that all four “corner” segments should display the same tolerance level. The data in Fig. 4.11 underline this principal symmetry, but we do see a slight discrepancy between the corner segments on the left and right side. We attribute this to a slight left-right misalignment of the LS with respect to the IrisAO. To demonstrate this, we proceed by measuring the individual contrast sensitivities of all segments (only the PASTIS matrix diagonal, with Eq. 4.15) for varying lateral misalignments of the LS. Starting from the nominal centered alignment, we move the LS by a fraction of the segment size, characterized by its circumscribed diameter of $D_{seg} = 1.4$ mm, as measured in the plane of the IrisAO. Then we run a couple of iterations of the WFS&C algorithm in order to optimize the DH solution and recover the nominal contrast level, measure the average contrast response in the DH to imposed (individual) segment pokes of 40 nm of WFE, and use Eq. 4.12 to calculate the WFE requirements per segment, for a target contrast of $\Delta c_t = 5 \times 10^{-7}$. Repeating this for five different misalignments along the x axis yields the data shown in Fig. 4.12, left. The nominal alignment case (dashed red) has been measured on a new DM solution for a DH contrast of $c_0 = 3.8 \times 10^{-8}$ (as opposed to the previous $c_0 = 2.5 \times 10^{-8}$), which is why the segment tolerances deviate slightly from the numbers indicated by the blue line in Fig. 4.11. While the tolerances for segments on top and bottom of the pupil (19 and 28) show no changes across the different misalignments, all other segments of the outer ring do. These results are directly connected to the exposed area of a segment in each alignment state of the LS: the more area is exposed the more stringent is its requirement. This is most visible for the four corner segments (22, 25, 31 and 34), as their area doubles between the farthest alignment states. The values for the side segments (23, 24 on the right and 32, 33 on the left) change as a function of misalignment too, but less so as their exposed area changes much less. The sensitivity analysis of single segments with respect to the position of the LS demonstrates the interest of the PASTIS tolerancing model: while the total WFE requirement, over the full pupil, is the same for all six LS states shown in Fig. 4.12, the segments display different individual tolerancing levels depending on where in the pupil they are located.

In order to validate the computed per-segment tolerances, we proceed by running a Monte Carlo experiment for both target cases shown in Fig. 4.11. For each experiment, we sequentially apply 1000 different WFE aberration patterns on the segmented DM and record the resulting average DH contrast values. The tolerances in Fig. 4.11 are the prescription determining how to draw these random WFE realizations: each segment-level WFE on segment k , in a single random WFE map \mathbf{a} , is drawn from its own zero-mean normal distribution with a standard deviation of μ_k , given by Eq. 4.12. This means that one random HiCAT WFE map is composed of 37 independent normal distributions with each a mean of zero, and a standard deviation of μ_k , which then gets applied to the IrisAO on HiCAT, and a DH contrast measurement is made. Since the tolerancing target is Δc_t , which is independent of the coronagraph floor c_0 , we intersperse measurements with a flat IrisAO in each iteration to capture the evolution of the contrast floor on the testbed. We subtract these $c_0(t)$ values from the DH measurements in the respective iteration in order to receive our final results in Fig. 4.13. We plot the time series of the experimentally measured contrast responses from the segmented WFE maps in the two left panels in Fig. 4.13. The green bottom curves depict the evolution of the contrast floor over time, with the IrisAO DM repeatedly being reset to its best flat. The top blue curves show the contrast measurements with random WFE map realizations applied to the segmented DM, as prescribed with the tolerances in Fig. 4.11, from which the green contrast floor has been subtracted. The histograms of the blue contrast curves are plotted in the two right panels in Fig. 4.13.

The resulting histograms have experimentally measured mean values of 5.12×10^{-7} and 3.04×10^{-7} (red dashed-dotted lines). These are very close to the target contrasts of 5×10^{-7} and 3×10^{-7} (orange dashed-dotted lines). This excellent fit between the target and experimental values ($\sim 1\text{--}3\%$) is clearly sufficient to prove our concept at the required level. The difference is larger for the standard deviations: the experimentally measured values of 1.3×10^{-7} and 7.9×10^{-8} (red dotted lines) are off from their analytically calculated values with Eq. 4.9 of 1.2×10^{-7} and 6.7×10^{-7} (orange dotted lines), by 5–7%. This is expected because of a slower convergence of variance estimators with respect to mean estimators. We note a slight asymmetry in the histograms, biased toward higher contrast. While the underlying assumption for the tolerancing method used in this paper is that the segment aberrations follow Gaussian statistics (Laginja et al. 2021, Sec. 4.1), which is reasonable for the aberrations on a segmented mirror telescope, the contrast itself does not follow a Gaussian statistic; it is the sum of squared Gaussian variables which is also known as a generalized χ^2 statistic, causing the asymmetry.

Simulated segment tolerances

Ideally, we could use a simulated model directly to be able to assess the segmented tolerancing levels for a particular contrast level. To this end, we use M^{sim} to calculate a set of segment tolerances for a target contrast of 10^{-6} and repeat the experiment described above, with 1000 iterations. Most efforts put into the simulator to date were intentionally invested in matching the operational interface, and the optical scales and morphology (e.g., orientations, sampling and location of diffraction features, photometry). The optical model currently matches its contrast predictions to results on the hardware to within a factor of a few, which is why in this experiment, we chose a less demanding target contrast compared to the experiments shown in Fig. 4.13. In this way, we retain an error margin for the resulting hardware contrast values that takes this into account. The results of this experiment are plotted in Fig. 4.14. We can see how the experimentally measured mean contrast (red dashed-dotted line) fails to meet the target contrast marked by the black dashed line with an error of 10%. This indicates that the tolerances derived from the simulated PASTIS matrix are less accurate than the experimentally measured ones.

Since the independent segment requirements in Eq. 4.12 (case of diagonal covariance matrix C_a) only depend on the PASTIS matrix diagonal, we compare the diagonals of the experimentally measured and simulated PASTIS matrix in Fig. 4.15. While the experimental matrix displays the higher absolute peak diagonal element (for the center segment), on average it contains lower contrast contributions per segment than the simulated matrix. In particular, the outer segments in the pupil are consistently over-estimated in terms of contrast contribution by the simulated matrix. A closer inspection of the pupil image shows that the segmented mirror model in the HiCAT simulator is in fact over-stretched along the x axis compared to the real IrisAO on the testbed, by 1.4%. As a consequence, the segments of the outer ring expose $\sim 5\%$ more total area compared to their visible area on the hardware, which in turn increases their influence on overall




Figure 4.13: Monte Carlo experiments on HiCAT to validate the independent segment error budget shown in Fig. 4.11, calculated from the experimental PASTIS matrix with Eq. 4.12, for target contrasts of $\Delta c_t = 5 \times 10^{-7}$ (top) and $\Delta c_t = 3 \times 10^{-7}$ (bottom). *Left:* Time series of the measured random contrasts (blue), from which the intermittently measured unaberrated contrast drift (green) has been subtracted. The total duration of each experiment is 65 min and captures 1000 randomly generated WFE maps. The target contrast of both experiments is indicated by the dashed orange lines. *Right:* Contrast measurement from the random WFEs to the left, plotted as histograms. The experimental distributions have a mean (dashed-dotted red lines) of 5.12×10^{-7} (top) and 3.04×10^{-7} (bottom). The expected standard deviations (dotted orange lines) as calculated by Eq. 4.9 are 1.2×10^{-7} (top) and 6.7×10^{-8} (bottom), versus the experimentally measured standard deviations (dotted red lines) of 1.3×10^{-7} and 7.7×10^{-8} .




Figure 4.14: Monte Carlo experiment on HiCAT for segment tolerances obtained from the simulated PASTIS matrix diagonal using Eq. 4.12. The target contrast $\Delta c_t = 10^{-6}$ as determined by the simulated tolerances is indicated with the dashed black line. The experimental mean over 1000 segmented WFE maps (dashed-dotted red line) does not recover the target contrast. However, we are able to predict this discrepancy analytically: using Eqs. 4.8 and 4.9, we can calculate the resulting contrast mean (dashed-dotted orange line) and variance (orange dotted lines) from the experimental PASTIS matrix and simulated segment tolerances directly. For details, see the main text.

contrast in the simulator. With an increased contrast sensitivity to these segments in simulation, it explains why the Monte Carlo experiment in Fig. 4.14 misses the targeted mean contrast: since the simulated matrix assumes a larger contrast influence by each individual (outer ring) segment, and the segment tolerances μ_k are inversely proportional to the matrix, the resulting per-segment requirements turn out more restrictive than they need to be in reality. This result shows that the simulated and experimentally measured PASTIS matrices have significant differences on their respective diagonals, which is not captured in Fig. 4.7. This is because the statistical analysis in Fig. 4.14 uses only the matrix diagonal, while the deterministic model in Fig. 4.7 (Eq. 4.4) uses the full PASTIS matrix. In the latter case, it turns out that the differences between M^{exp} and M^{sim} average out.

While the discrepancies between the experimental and simulated matrix lead to an offset in the tolerancing results that can be improved upon with a more accurate model, we can show that this offset is directly predictable by using the experimentally measured matrix M^{exp} with Eqs. 4.8 and 4.9. We consider the segment tolerances obtained with M^{sim} to be a covariance matrix C_a^{sim} , with the segment variances filling the covariance matrix diagonal. This allows us to calculate the analytically predicted mean with $\text{tr}(M^{\text{exp}} C_a^{\text{sim}}) = 8.75 \times 10^{-7}$ (dashed-dotted orange line), and the variance with $2 \text{tr}[(M^{\text{exp}} C_a^{\text{sim}})^2]$, with its square root yielding a standard deviation on the contrast of 2.2×10^{-7} (dotted orange lines). These values accord with the experimentally measured mean of 9.07×10^{-7} (dashed-dotted red line) and standard deviation of 2.3×10^{-7} (dotted red lines) to within a statistical error. Usage of these formulae circumvents the need to reevaluate the contrast response of a large number of segmented aberration maps when a new segment covariance matrix C_a is obtained from modeling, and can be used for the direct assessment of the contrast performance for by-design instruments.

Overall, our experiments on HiCAT present successful experimental validations of the PASTIS model for a specific high-contrast instrument. We measured an experimental PASTIS matrix and validated it by




Figure 4.15: Comparing the PASTIS matrix diagonals from the hardware matrix M^{exp} (dashed orange line) and simulated matrix M^{sim} (solid blue line). Segment number 0 (center segment) in the experimental matrix has the largest absolute contrast influence, but the simulated matrix assumes on average the larger influence of the individual segments on the average DH contrast. This translates directly into more stringent segment tolerances, μ_k , than necessary for the given hardware setup.

comparing its modeled contrast results to testbed measurements. We decomposed the matrix into independent optical modes that we scaled uniformly and cumulatively to a differential target contrast of 10^{-6} . We calculated statistical segment-level WFE tolerances under the assumption of independent segments and validated them with Monte Carlo experiments at target contrasts of 5×10^{-7} and 3×10^{-7} , measuring the contrast from randomly drawn segmented WFE maps as prescribed by the derived requirements. While using tolerances derived from a simulated PASTIS matrix was not enough to reach a particular contrast goal, the experimentally measured matrix allowed us to validate the analytical contrast predictions in terms of a contrast mean and variance for an arbitrary segment covariance matrix. Future work will aim to optimize the differing segment sizes in the segmented DM model; however, matching the contrast prediction of the HiCAT simulator to the hardware is out of the scope of this study.

4.5 Discussion

The astronomical community's experience with space-based segmented observatories is currently limited to JWST, which will be launched very soon. Because of gravity and thermal constraints, the telescope was not aligned entirely on the ground to test its optical performance. Instead, a large number of ground optical tests including interferometry and other metrology techniques were performed to validate the observatory-level optical model, including for example radius of curvature, or interferometric alignment of adjacent segments (Perrin et al. 2018; Knight et al. 2012b). Something like LUVOIR would be two to three times larger than JWST and therefore it is anticipated that performance validations of the observatory will rely heavily on model-based assessments due to the sheer size and complexity of the telescope.

This raises the necessity for modeling, performance prediction and tolerancing tools that are not dependent on testing of the fully integrated observatory system. The PASTIS tolerancing tool, presented in theory so far (Laginja et al. 2021; Leboulleux et al. 2018b), provides us with methods to derive segment stability tolerances that in turn will determine the coronagraphic performance of an imaging instrument like those on LUVOIR. In this paper, we perform the first hardware validation of an experimentally calibrated tolerancing model. We use this to determine the constraints on segmented mirror stability for various




Figure 4.16: Comparing independent segment WFE tolerances for different levels of differential target contrasts, calculated from the experimental PASTIS matrix M^{exp} . The tolerances for $\Delta c_t = 5 \times 10^{-7}$ and $\Delta c_t = 3 \times 10^{-7}$ have been validated experimentally in the current performance regime on HiCAT. The other lines describe predicted requirement limits on the segmented mirror for deeper contrast levels. These curves isolate the influence of segmented mirror drifts from other system components that influence the contrast. The curves appear top to bottom in the same sequence as in the figure legend.

levels of contrast on the HiCAT testbed, which can be used to derive requirements of future testbeds that will perform system analyses for the LUVUOIR mission. This paper validates the statistics of phasing errors without making any timescale assumptions for the WFE. Therefore, this applies not only to static phasing residuals (which can inform the design of a sensing and control strategy), but also to dynamic tolerances in the context of adaptive optics (AO) loops.

During the experiments, we observed that the variability in the static contrast poses a problem to the accurate measurement of the experimental PASTIS matrix for the mirror segments, and the tolerancing validations. We solved this issue in good part by adopting the differential contrast Δc as the metric of interest: We have introduced a reformulation of the PASTIS tolerancing formalism in order to separate the contrast influence of segmented mirror misalignments from all other aberration sources in the instrument (see Sec. 4.2). This is particularly useful when we assume a time-dependent phase aberration term $\phi_{ab}(\mathbf{r}, t)$ which will influence the absolute raw contrast.

We successfully validated the PASTIS model on a real high-contrast instrument by measuring an experimental PASTIS matrix M^{exp} . We performed the individual segment tolerancing with this matrix and have shown that the derived segment requirements are indeed the correct standard deviations for a targeted mean contrast. The experimental results in Fig. 4.13 coincide very well with the analytical formulas for contrast mean and variance in Eqs. 4.8 and 4.9, down to a statistical error.

We can use this method to extrapolate the stability requirement of the segmented mirror to contrast levels that lie beyond the current performance capability of the testbed. In the case of HiCAT, we are currently limited to a static contrast of 2×10^{-8} , and contrast variations up to the order of 1×10^{-7} due to influences from the testbed environment and an incoherent background attributed to the light source. After validating the individual segment requirements in a contrast regime in which the contrast floor and drift do not pose any problems, we compare the individual segment tolerances as calculated with M^{exp} in Eq. 4.12 for various target contrasts down to the 10^{-8} regime in Fig. 4.16. This result informs us about the level of the segmented mirror contrast stability for future demonstrations that intend to get closer to the envisioned LUVUOIR performance. For example, the segmented mirror needs to be able to keep a sub-nanometer stability if we want to achieve a contrast of 10^{-8} . The numbers in Fig. 4.16 exclusively show the requirements on the WFE component caused by the segmented mirror, which builds a fundamental piece in the total error budget of a segmented high-contrast instrument.

The mean differential contrast $\langle \Delta c \rangle$ can be written as a sum of unrelated contributions from the segments only if the segment pistons are statistically independent (see Eq. 4.11), that is, if the matrix C_a is diagonal. But even if the segment pistons are correlated and even if these correlations are unknown, the mean differential contrast *can* be written as a sum of independent contributions, namely those of the optical modes of the system (see Eq. 4.10).

We showed that it is possible to use Eq. 4.17 in order to scale the individual optical modes to yield a particular target contrast (Fig. 4.10). While there remain deviations from the exact equal contrast contribution of each mode to the overall contrast, this result clearly demonstrates that we can impose a weighted scaling on these modes in order to influence the contrast. The remaining errors we see in the uniform mode tolerancing are likely coming from calibration errors in the segmented DM actuator influences, unidentified aberrations in the system, and an insufficiently calibrated contrast floor for the short duration of this experiment (~ 5 min).

Fundamentally, these optical PASTIS modes represent the natural instrument modes with respect to contrast, which can be exploited for closed-loop adaptive optics (AO control during the operation of the instrument). The LUVOIR mission aims to deploy a full AO system in space, which requires control on all spatial frequencies. While fast, low-spatial-frequency control can be done with a low-order wavefront sensor (LOWFS), a high-order wavefront sensor (HOWFS), which is not subjected to spatial filtering, will be sensing aberrations on the segmented pupil (Pueyo et al. 2021; Pogorelyuk et al. 2021).

Optimal wavefront control strategies will need to optimize the DH contrast, which is the main science metric for these future instruments. The PASTIS modes therefore offer a natural application to this problem since their contrast influence is directly quantified by their respective eigenvalues. PASTIS modes should therefore be investigated further as part of the design of the high-order modal control scheme of a multiple-layer space AO system for high-contrast applications.

There are existing methods (Chambouleyron et al. 2021) that allow a Fourier-filtering WFS to be designed with given sensitivity to specific modes. Exploiting this, one could design a WFS FPM, and thus a transmissive mask, giving an enhanced sensitivity to the modes degrading the contrast (whose influence mainly falls inside the DH area), accepting to have a reduced sensitivity outside. With such a WFS, the control loop could be optimized to maximize the contrast performance of the instrument directly.

4.6 Conclusions

Accurate tolerancing of different WFE contributions on future large observatories is crucial in order to be able to design systems capable of sufficiently stable contrast levels for exoEarth detection, to predict their contrast stability, and to assess their performance. In particular, the WFE contributions from cophasing errors on segmented telescopes will have a direct impact on the performance of the high-contrast instrument. In this paper, we used the PASTIS tolerancing model to perform a segmented WFE tolerancing analysis on the HiCAT testbed, and presented experimental validations to demonstrate its utility.

We successfully measured an experimental PASTIS matrix on a 37-segment IrisAO mirror after isolating the influence of the segments in the overall contribution to contrast drift. The individual segment tolerances calculated from this matrix yield an accurate mean contrast and variance in Monte Carlo experiments when compared to analytical predictions, up to a minimal statistical error. We also compared these experimentally obtained segment tolerances to equivalent results obtained from a simulated PASTIS matrix. The experimental measurements were more accurate for performance predictions, but the errors from the simulated segment tolerances can likely be minimized with a more accurate model of the segmented mirror.

Combining the experimental PASTIS matrix, which represents the realistic contrast influence of the testbed, with a covariance matrix that describes segment piston variations allowed us to correctly predict the resulting contrast mean and variance. This allows for a simple evaluation of the expected contrast stability of a per-design instrument. Such a covariance matrix can be a diagonal one to describe the independent segments from a simple segmented mirror design, as in the HiCAT case, or a non-diagonal one that incorporates knowledge from opto-mechanical correlations between segments coming from realistic finite-element modeling.

We use the experimentally measured matrix to predict the required wavefront stability of the segmented DM on HiCAT for contrast levels that are currently out of reach due to environmental influences on the testbed. We first validated the segment tolerances for a differential contrast of 5×10^{-7} and 3×10^{-7} , for which the segment requirement standard deviations range from 2.5 to 6 nm. We then proceeded by establishing a set of requirements for the segmented mirror for a contrast contribution in the 10^{-8} regime, and conclude that the wavefront stability from the segmented DM will have to be better than 1 nm for each individual segment.

Our future work aims to measure a simulated PASTIS matrix with a more accurate model and use it to derive WFE tolerances that correctly define the stability requirements for a target mean contrast measured on the hardware. Further, we aim to demonstrate how to measure a PASTIS matrix on one contrast level, and use the derived tolerance limits on a better performing contrast level. Finally, we intend to explore closed-loop modal control with a HOWFS by using the optical modes from a measured PASTIS matrix, since these modes represent the direct sensitivity of the instrument contrast to segmented mirror misalignments.

5

APPLICATIONS OF ANALYTICAL TOLERANCING IN STUDIES FOR ULTRA-STABLE TELESCOPES

My work on segmented coronagraph sensitivity analyses started as part of the ULTRA study contract between Ball Aerospace and our group at STScI. Lucie Leboulleux had developed the PASTIS forward modeling approach during her PhD to serve as an alternative for numerical end-to-end simulations in tolerancing studies for large, segmented telescopes. It is only after a closer look was put on the new statistical interpretation and inversion of the model that we realized how useful it would be in these studies. The work I presented in Chap. 2 lays out a very fundamental framework for the statistical tolerancing of aberrations through a coronagraph, and it enabled a wide range of applications. It has been used by other researchers, both internal and external to our team, during my PhD already to continue the efforts of learning how to define allowable WFE limits on future HCI missions.

In a wider context, I broadened the code and method I developed during my initial work to enable a full sensitivity analysis on an extended range of aberration modes beyond localized WFE on segments. This allows us to treat the full spatial frequency range of expected aberrations on a large, segmented telescope. Already discussed in the paper of Chap. 2, I mention that the PASTIS tolerancing framework can be used for thermo-mechanical constraints on telescope structures. Indeed, I continued my work on analytical tolerancing by using models of thermal and mechanical responses of individual segments, provided by L3 Harris Technologies, to perform a quantitative analysis of the physical requirements for such modes. I present this joint study with Laurent Pueyo and Ananya Sahoo, and its results, which are currently a journal paper in preparation that I am a co-author on, in this chapter.

A central point of concern for the definition of WFE requirements is the amplitude of allowable WFE drifts. There is a range of typical timescales associated with the various spatial frequency components of WFEs, which need to be considered when creating error budgets. While the static control approach traditionally deployed on space telescopes could in theory be extrapolated to the anticipated large structures of an exoplanet-hunting space observatory, this results in dauntingly tight WFE requirements. Especially segmented primary mirrors exhibit very stringent requirements of only a couple of picometers across the entire multi-meter structure in this case, maintained over the course of several hours. I have already shown in Chap. 2 that this is too conservative, as not all segments in the pupil need to be held equally stable. Moreover, it is here where space-based telescopes benefit tremendously from the knowledge developed by the ground-based community, which offers enormous expertise in adaptive optics and associated WFS&C methods for closed-loop wavefront control. In the very early stages of my PhD, I spent many hours on a white board to understand the analytical connections between coronagraphic sensitivities, wavefront sensing efficiency and WFE drift rates that Laurent Pueyo had devised, which was substantial to formulate a justification for the PASTIS tolerancing model I published as my first PhD result. I use this chapter to introduce the analytical framework that allows us to put my results for segment tolerancing onto appropriate timescales for dynamical analyses. The major result is: Folding the segment-by-segment tolerancing maps we obtain through the PASTIS analysis method into a paradigm that allows for continuous correction of the drifting aberrations significantly relaxes the imposed WFE requirements.

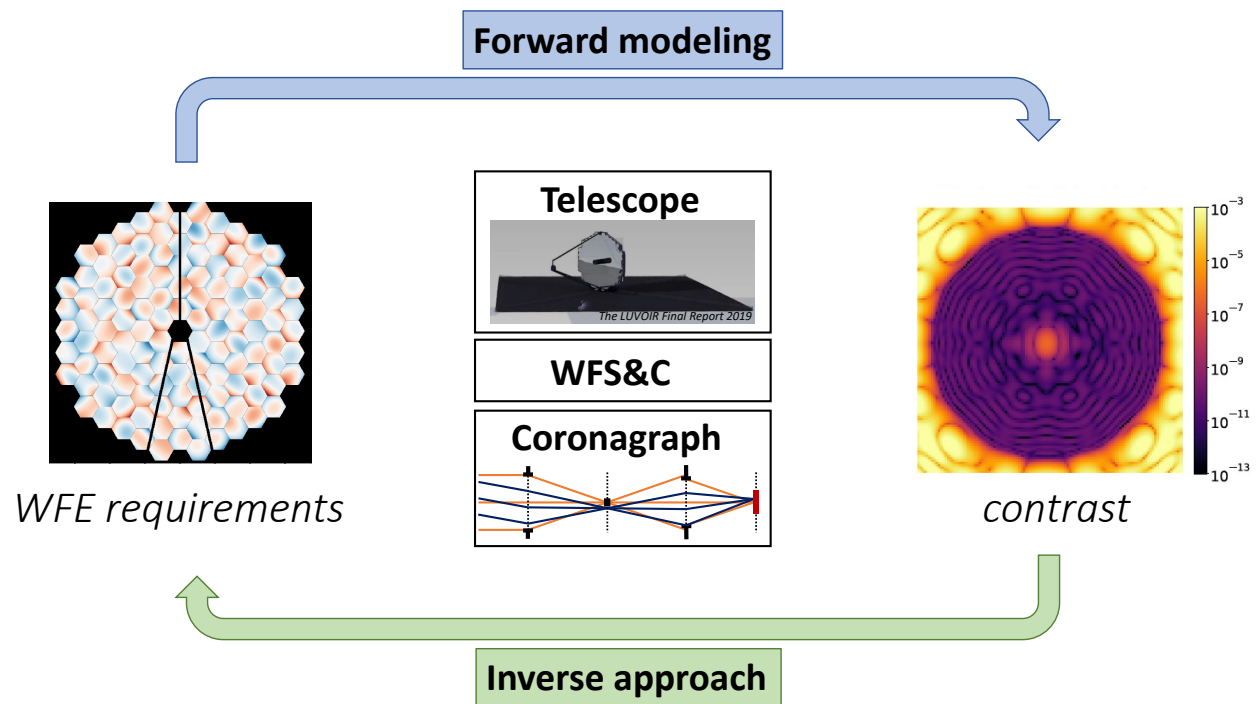


Figure 5.1: Exploring the relationship between telescope, WFS&C strategy and coronagraph is the main goal of WFE sensitivity analyses. Forward modeling is the more traditional approach in which we numerically propagate different states of the instrument to predict the HCI performance. Inverse modeling, the approach followed in this thesis, starts from the required imaging performance and derives WFE requirements from that. This method necessitates proper analytical models of the optical system and its control architecture.

5.1 Tolerancing approaches for segmented telescopes

In the big picture, what we are exploring is the relationship between the telescope and all its subsystems, the WFS&C strategy and the coronagraph, as illustrated in Fig. 5.1. It is crucial in this context to design, construct and control the telescope and coronagraph as an integrated system and across a range of spatial and temporal scales, which makes the definition of detailed subsystem WFE requirements no trivial task. The traditional way to approach this issue is forward modeling. In this case, some assumptions about the telescope are being made, aberrations with a certain WFE are added to the optics, and end-to-end simulations then predict the performance with a given WFS&C strategy and coronagraph architecture (Potier et al. 2021; Juanola-Parramon et al. 2019a,b). This process is repeated a large number of times, usually in the form of Monte Carlo simulations, in order to converge to appropriate WFE specifications for the targeted contrast by a trial-and-error approach. While this allows us to constrain the tolerances to a certain degree, there is no clear procedure in which the contrast level is the input and the aberration limits are the output. Moreover, this does not permit any conclusion about individual segment tolerances, as this would increase the dimension of the parameter search to an untractable level.

The alternative approach is to set science requirements, for example motivated by exoplanet yield, and then invert the aberrations-to-contrast relationship, at the heart of which lies the system consisting of the telescope, WFS&C system, and coronagraph, and compute the WFE limitations from that. Optical simulations can be used to convert contrast to wavefront stability, taking the various spatial and temporal scales into account, and there has been work that uses such analytical approaches to quantify stability requirements this way (Nemati et al. 2020; Stahl et al. 2020; Nemati et al. 2017a).

While the forward modeling approach predicts the imaging performance, the inverse approach allows us to first define top-level contrast stability constraints, which are fixed by the coronagraph performance and expected science yield, and then compute WFE requirements as a direct result. While they gives us an answer for the entire system as a whole, the next point to be addressed is how to flow these top-

level requirements down into all the sub-systems of the observatory - where do we allocate how much of the total error budget? The recent program for Ultra-stable Large Telescope Research and Analysis (ULTRA; Coyle et al. 2019b) led by Ball Aerospace performed a system study to address technologies that are needed for picometer-level optical stability on large telescope structures, and addressed in particular stability budgets. The classical error allocation used in these cases is a straightforward root-sum-squared (RSS) (Lightsey et al. 2014), in which the overall WFE requirements get distributed to relevant subsystems. Within the scope of studies like ULTRA however, it was found that this might be overly constraining in some components of the system, and a approach based on power spectral densities (PSDs) to rebalance the allocations in order to avoid an excessive strain on any one element or technology would be more appropriate.

The tolerancing approach with the PASTIS model finds its place at the very basis of these relations. Its method of deriving WFE requirements offers itself exemplarily for the treatment of the mid-spatial-frequency regime consisting of segment-to-segment aberrations, but the analytical nature of the work is readily applicable to the treatment of any arbitrary aberration basis, which is what I dedicate the first part of this chapter to. This provides us with a unified way of treating coronagraphic sensitivities with respect to the full range of spatial scales on which WFEs are to be expected, and leads to some very useful results in the physical tolerancing of thermal gradients across segments.

One of the major results emerging from the work in Chap. 2 is the fact that the segments of a large telescope do not exhibit the same sensitivity to contrast across the aperture. Instead, there is a spatial structure in their allowable WFEs, which supports the idea that the tolerancing allocation in separate subsystems of the HCI instrument need to be done very carefully. Mostly so, because it illustrates that the often-quoted picometer-level WFE stability only concerns a small subset of possible misalignment modes, in certain parts of the pupil. In turn, the segments and aberration modes that are not touched by this worst-case-scenario require stability levels much closer to what we have already been able to achieve with the nanometer-level requirements on for example JWST.

Considering the statistical approach to tolerancing that the PASTIS method brings with it, as we described at length in Chap. 2, it is important to point out that PASTIS itself makes no assumptions about the timescales involved in the tolerancing problem, as it does not assume any temporal evolution. The framework itself can thus be used directly for the definition of static WFE limitations for a requested contrast level, which is a crucial aspect in the design of passive telescope stability and control strategy. Equally, it can be used as an input to dynamic studies, pairing it with the appropriate aberration timescales of concern in a large telescope structure. The second part of this chapter dives into this dynamic WFE tolerancing, and treats the bigger picture of defining dynamic WFE on a large, segmented telescope.

The idea is to have a proper analytical model for the propagation of aberrations, given by PASTIS, which can be folded into the paradigm of operating a full AO system in space, as mentioned in Chap. 1.4.2. PASTIS establishes an analytical relationship between the change of contrast, the linear coronagraph model and the wavefront covariance, whether static and open-loop or dynamic and closed-loop, as its conclusions are purely statistical. Therefore, it provides the ideal basis for the definition of WFE drift requirements, when paired with the pertinent timescales of a WFS&C architecture.

In Sec. 5.2, I present the extension to modal bases of all spatial frequencies. In Sec. 5.3, I first detail the global ideas behind dynamic WFE tolerancing, and the analytical equations used for it. I proceed by showing how to apply this to dynamic segment-level tolerancing under the inclusion of an ideal wavefront sensor, and show the dynamical tolerancing maps for LUVUOIR A that are a direct result of my work with PASTIS.

5.2 Extended spatial tolerancing bases

Since large space telescopes typically have various options for wavefront control (WFC), addressing different sources and spatial frequencies of the aberrations, it helps to categorize aberrations into three spatial frequency domains: low, mid and high spatial frequencies. Such a separation between three spatial frequency bands originates mostly in the availability of mitigation strategies for WFEs in the separate regimes, an approach that has been identified useful for JWST (Lightsey et al. 2014), and adapted later for initial




Figure 5.2: Pupil-plane aberrations with various spatial frequencies (top) and their response in the coronagraphic dark hole (bottom).

error budgets for LUVOIR (Lightsey et al. 2018). For example, using movements of the secondary mirror is considered more amenable to control global aberrations like defocus or astigmatism, while segment-level aberrations are addressed by adjusting the degrees of freedom on individual segments. While the chosen control strategy represents one defining factor for the different spatial frequency regimes, another one is the optical response in the image plane. Due to the Fourier optics propagation of aberrations to the focal plane, the impact of an aberration on the DH will depend on its spatial frequency content. The distinct regimes thus have an impact in different parts of the DH, as illustrated in Fig. 5.2: low-spatial-frequency aberrations will induce light close around the optical axis, high frequencies will direct light outside the dark hole and to concentrated specks of light, while mid frequencies cause a response in the middle of the DH, which is the most interesting area in the focal plane for exoplanet observations.

Low-spatial-frequency aberration modes stem from global misalignments of the telescope, in particular secondary mirror misalignments and global deformations in the backplane structure. The aberrations are easily described with a Zernike modal basis (Noll 1976) and can be sensed with a low-order wavefront sensor. These aberrations can be compensated with motions of the secondary mirror, but many coronagraphs are intrinsically insensitive to these low-order modes (N’Diaye et al. 2015b; Green & Shaklan 2003) due to their robustness against stellar angular size, baked into their design. Mid-spatial-frequency aberrations arise from segmented aberrations on the primary (Douglas et al. 2019; Ruane et al. 2017), which means that they do not exist on monolithic telescopes. They can be sensed with edge sensors (Coyle et al. 2018; Lou et al. 2018) and an out-of-band WFS that is not spatially filtered by the coronagraph (Moore & Redding 2018). While these WFEs are often mitigated with direct control of the segments themselves, they can also be addressed with the main control components of the WFS&C loop, namely the continuous DMs, or offloaded to a segmented DM in a relay pupil (Lumbres et al. 2018). Finding the ideal modal basis to describe the impact of these mid-frequency aberration modes on the contrast was presented in Chap. 2, and in this chapter, I show how to analyze such aberrations beyond just a local piston misalignment. High-spatial-frequency aberration modes arise mostly from surface polishing errors, which avails itself straightforwardly to a treatment with sinusoidal functions, that is “ripples” across the pupil.

The PASTIS model provides a fundamental tolerancing framework for any set of input aberrations. The formalism is applicable to any arbitrary modal basis of WFEs. In the following sections, I present the application to the full range of spatial frequencies as described above. Nominally, the local and global




Figure 5.3: Examples of global aberration modes as an alternative aberration base for tolerancing analyses with PASTIS, in form of their phase on the primary mirror. The top row shows the Zernike modes tilt, defocus and vertical quadrifoil, each with a surface error amplitude of 1 nm rms. The sine wave aberrations on the bottom have different spatial frequencies and orientations, and all have a surface error amplitude of 10 pm.

modes as described by these three spatial frequency regimes need to share a contrast allocation for a particular target contrast. However, as they are all addressed with different control strategies as described above, it is easier to treat them separately, as if the modes of a chosen basis (only local modes, or only low-order global modes) were the only source of contrast degradation. Performing a comprehensive analysis of all of them together can then be achieved by adjusting their individual contributions to the target contrast.

I briefly address the application of low- and high-spatial frequency aberrations over the entire telescope pupil, before extending the tolerancing application to local Zernike modes, and to thermally and mechanically modeled local aberration modes by L3 Harris Technologies.

5.2.1 Global modes for low and high spatial frequencies

One of the original motivation points for the implementation of semi-analytical PASTIS matrices that are calculated from a numerical simulator, as demonstrated in Luginja et al. (2019), was to expand the tolerancing application to non-analytical modes beyond what could be done with the original, fully analytical approach in Leboulleux et al. (2018b). In this section, we briefly touch on how to execute a sensitivity study for global aberration modes. This can be done for example by using the simulator described in Chap. 3.1 to impose global Zernikes (low spatial frequencies) and sinusoidal ripples (high spatial frequencies) on the telescope aperture, as shown in Fig. 5.3. Creating a contrast sensitivity matrix for global aberration modes means that the contrast is measured for two such modes applied simultaneously, and each row and column of the resulting matrix corresponds to one mode. For sinusoidal ripples for example, each row or column corresponds to a different spatial frequency and orientation across the telescope pupil. More generally, the PASTIS matrix dimensions now represent each individual mode of its modal basis (which can be per-segment, if so chosen), rather than representing an aberration per segment by default. Performing the analytical inversion of the PASTIS matrix for purposes of tolerancing can be done in any of these cases, following the formalism outlined in Chap. 2.3. In the general case, the analysis then yields WFE requirements per each mode of the used basis.

Global modes like these have been extensively studied, as the treatment of both low-order (Sidick & Riggs 2019; Riggs et al. 2019; N'Diaye et al. 2015b; Shaklan & Green 2005; Green & Shaklan 2003) as well as high-order spatial aberration modes, as they appear in any telescope, monolithic or segmented. For the

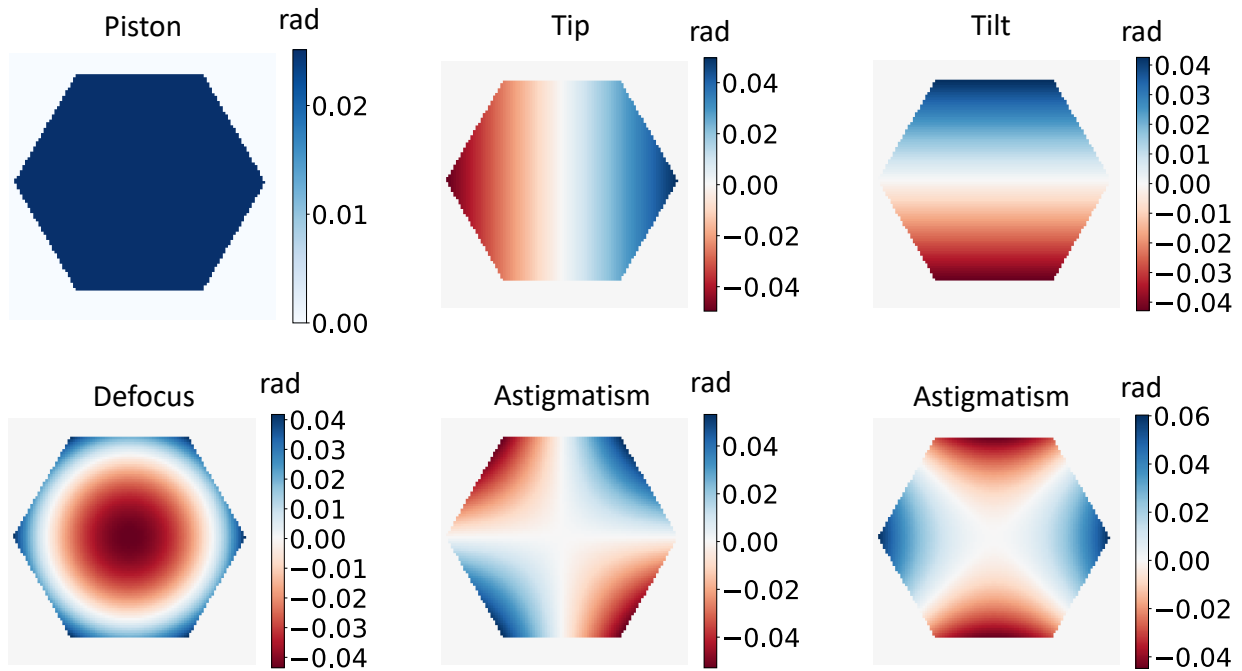


Figure 5.4: Examples of the first six local Zernike modes as ordered by Noll, approximated on a hexagonal support. Just as the pure piston mode used in the previous chapters of this thesis, they can be used for local segment tolerancing analyses.

purpose of this thesis, efforts were put in the formalization of the segment-level contributions, as explained in the following sections.

5.2.2 Local Zernike modes

The process for evaluating segmented tolerances with respect to contrast sensitivity is easily extendable to other modes than piston, as it follows the same procedure like shown in the previous chapters. The most straightforward modal basis for segment-level aberrations are Zernike polynomials (Noll 1976) truncated to hexagonal segments, as shown in Fig. 5.4. Any of these modes can be used in the formalism laid out in Chap. 2 to calculate a PASTIS matrix.

The same thing can be done on several local modes at once in order to obtain sensitivity maps that show the relative contrast influence between a whole set of aberration modes. Such a PASTIS matrix collects the relative contrast influence of the total number of used local modes, n_{loc} , on all segments, n_{seg} . The dimensions of such a matrix are thus $n_{loc}n_{seg} \times n_{loc}n_{seg}$, and the mean target contrast is attained by combining all toleranced modes together. The resulting tolerance maps of such a process, for the first six Zernikes, are shown in Fig. 5.5, performed on the small-angle APLC of the LUVUOIR A observatory, equivalently to Chap. 2.5. This multi-mode tolerancing approach can be readily adapted to an arbitrary modal basis, for example the custom local modes from the following section.

5.2.3 Custom local modes from thermo-mechanical modeling

Zernikes are an idealized basis and while some of these modes have been shown to have a significant influence on the coronagraphic contrast (Juanola-Parramon et al. 2019b), they are unlikely to materialize as a segment deformation in this exact modal basis. Instead, we can perform a tolerancing analysis with proper segment modes from thermal and mechanical modeling provided by L3 Harris Technologies that describe the local segment behavior on a real telescope. These modes describe a segment's response as a function of applied thermal gradients or mechanical influences. A subset of thermal modes is shown in Fig. 5.6. These modes show the surface deformation of the mirror substrate generated by applying a temperature change of 1 milli-Kelvin (mK) along various axes of a segment. The mode in the center is the resulting deformation when heating the front end of a segment, the mode on the bottom right shows

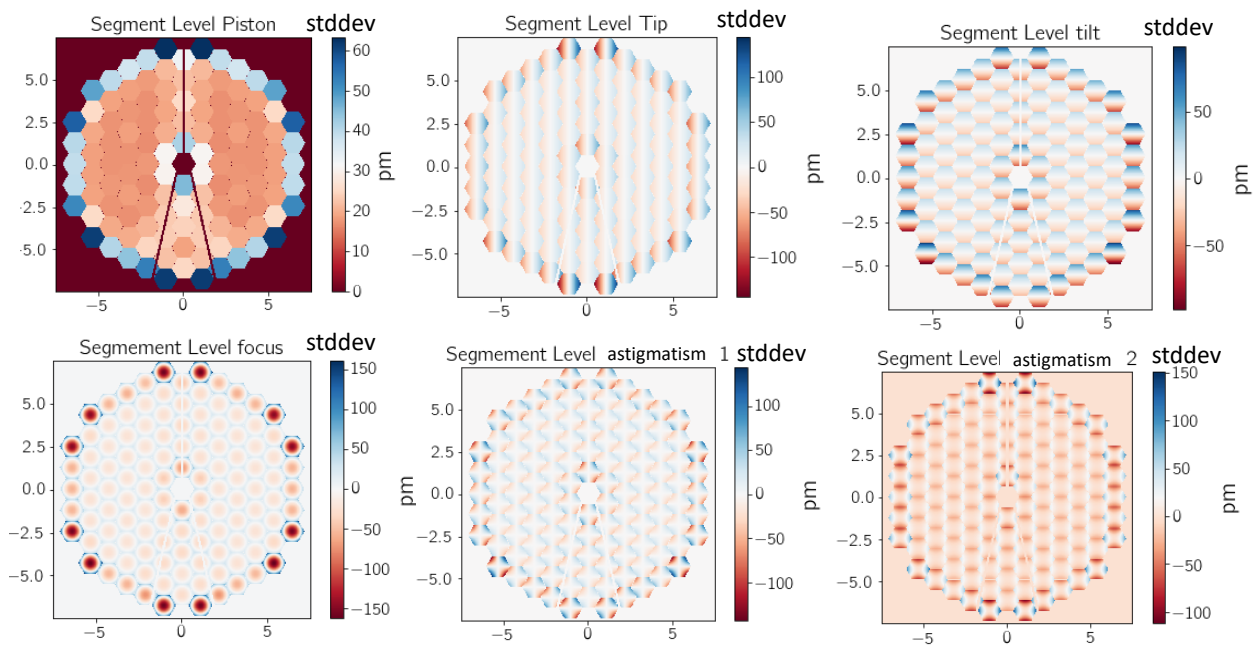


Figure 5.5: Segment-level requirements for a set of six local Zernike modes on LUVOIR with the small-design APLC, for a target contrast of $\Delta c_t = 10^{-10}$. A multi-mode analysis approach like this provides a way to do WFE tolerancing within a large number of aberration modes at once, taking into account requirement trade-offs between them as the requirement allocation is done for a uniform contrast across all segment modes.




Figure 5.6: Examples of thermo-mechanically modeled segment modes by L3 Harris Technologies. They show the surface deformation for a temperature change of 1 mK applied in different ways, and along various spatial axes. This data allows us to connect changes in WFE with physical tolerances, in this case temperature gradients.

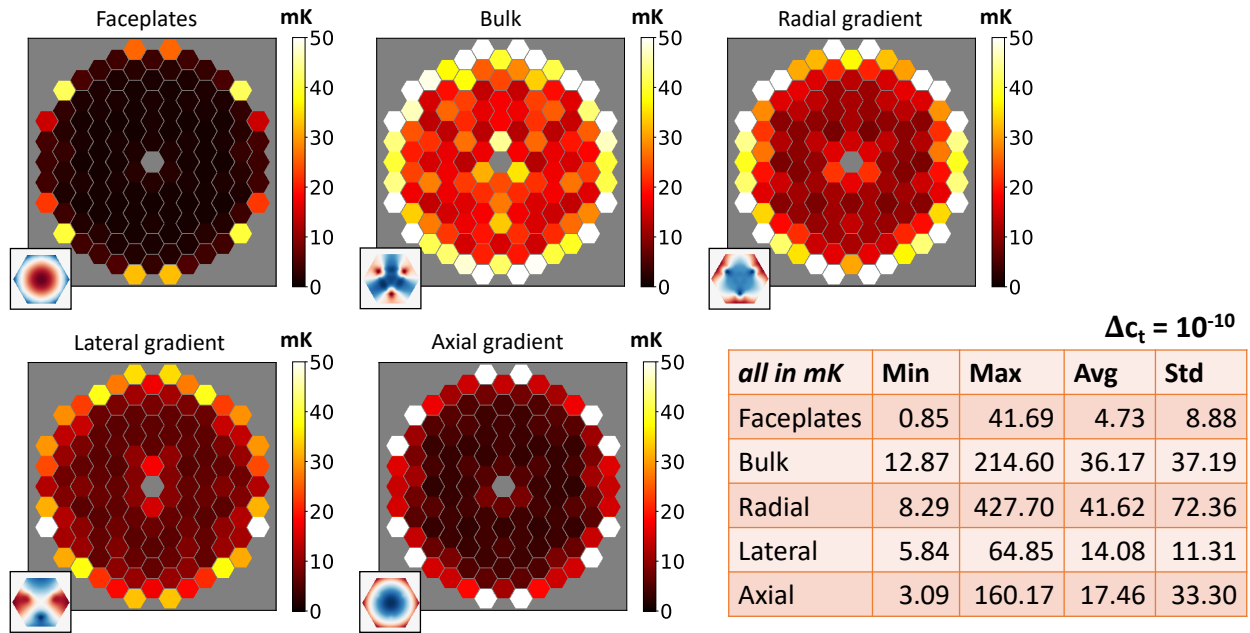


Figure 5.7: Acceptable thermal gradients are a direct physical interpretation of the surface requirements as computed with PASTIS, interpreted as a standard deviation in mK. These are shown here for the requirements of the five thermal Harris modes, where all of them together, on each segment, are tolerated to a total target contrast of $\Delta c_t = 10^{-10}$, on the narrow-angle APLC of LUVUOIR A. The inset figures below each requirement map indicate the used local mode from the surface plots shown in Fig. 5.6.

the shape change when the whole segment is heated in bulk. The remaining three maps display a 1 mK gradient along the radial, x-lateral and axial (z) axes.

We can use these local modes as a basis to do segment-level error budgeting, that is to quantitatively constrain the surface deformation of each segment given an overall target contrast. By using thermally modeled modes, this allows us to do the tolerancing in physical units. The deformation models give us the relationship between a thermal gradient and the surface deformation, and the PASTIS tolerancing model establishes the connection between surface deformation and contrast degradation. This means that we can use the PASTIS framework to determine thermal tolerances motivated by wanting to limit the coronagraphic contrast degradation.

We create a PASTIS matrix by applying these modes sequentially to each segment on the narrow-angle APLC on LUVUOIR A, and calculate sensitivity maps for each of the local, thermal modes shown in Fig. 5.6. As in Chap. 2.4.2, we assume that all basis modes are independent of each other and that we have access to the statistical mean of the average contrast in the DH. With five individual local modes applied to each of the 120 segments in the LUVUOIR pupil, the total number of modes used as a basis for this sensitivity analysis is $n_{modes} = n_{loc} \times n_{seg} = 5 \times 12 = 600$, which means that we are building a 600×600 PASTIS matrix.

For a total target contrast of $\Delta c = 10^{-10}$, we first calculate the per-segment tolerance maps for the Harris modes in terms of optical deformation (in pm) with PASTIS. Since these are expressed as one standard deviation per local mode and segment, we can convert them into a temperature scale in order to limit the allowable temperature change per segment. We rescale the tolerances from pm to mK with the respective thermal modes and plot the results in Fig. 5.7. Each segment now has a temperature value associated with it that corresponds to applicable temperature gradient limitations, expressed as a standard deviation. As with the optical tolerances from Zernike modes on this APLC, the inner segments need to be constrained more tightly than the outer segments, by 2–3 orders of magnitude. While the outer segments are more tolerant to deformations overall, the different modes still have a different effect on the various segments of the outer ring. The surface mode from a lateral gradient in particular (bottom left in Fig. 5.7) seems to influence the contrast more uniformly across the outer ring segments than for example the axial gradient modes (bottom center), which influence the most geometrically exposed “corner” segments in particular.

Associating a single physical value to each segment, it becomes more obvious which local modes have a stronger impact on the contrast than others. With all requirement maps in Fig. 5.7 being shown on the same temperature scale, we can see that on top of the relaxation of the outer ring across all modes, some modes influence the contrast globally less than others. Bulk heating of each individual segment (top center), for example, impacts the contrast significantly less than a temperature gradient imposed by faceplate heating (top left).

Using a specifically modeled local basis like this allows us to investigate aberration modes that have been identified through thermo-mechanical modeling to occur in a real telescope. Performing the full, statistical tolerancing analysis on this sort of data can translate the computed WFE requirements into physical units. This reformulates the initially defined science requirements as concrete thermal and mechanical limits that have to be met with the telescope design, plus the WFS&C and the real-time thermal control, in order to stay within the needed specifications, as set by the top-level mission goals.

5.3 Temporal domain analysis

While segmented telescopes with diameters over ten meters promise the most fruitful science results in terms of exoplanet yield, they are also more structurally complex and harder to keep stable than smaller telescopes, especially if monolithic. Conceiving such a large, segmented observatory as a space-borne extreme AO system can significantly relax the resulting WFE requirements, which is, incidentally, supported by the increased sensitivity tied to its large collecting area. Continuous WFS&C introduces new instrumental parameters that can be exploited in order to keep the contrast stable. We are now trading between the coronagraphic robustness to aberrations with the sensitivity of a wavefront sensor, which become the defining parameters when establishing allowable WFE drifts. We first go over the general analytical framework for WFE drift tolerancing in Sec. 5.3.1, and compare the relevant variables between an open-loop scenario, also known as a “set-and-forget” control strategy, with a closed-loop setting using continuous WFS&C. Using this generalized approach, partially presented in Pueyo et al. (2019) and Pueyo et al. (2021), we apply it to the open-loop results of individual segment tolerancing obtained in Chap. 2 to derive dynamic segment tolerancing maps in Sec. 5.3.2, and we discuss how this fits into the global scope of preparing a LUVUOIR-type mission.

5.3.1 Dynamic wavefront errors and relevant timescales

In this section, we compare the open-loop (OL) and closed-loop (CL) wavefront control strategies and identify relevant timescales for each. We further show an analytical derivation for the calculation of required dynamic WFE drift rates for both scenarios, and identify the key instrumental parameters that influence them. A visual representation of the contrast evolution in OL and CL operations is shown in Fig. 5.8. To perform a high-contrast observation on a large telescope from space, the first thing to do is to dig an initial DH by deploying a DH algorithm to reach the target raw contrast c_0 , at an initial point in time t_0 , which includes a compensation of static WFEs. It is envisioned to do this “DH digging” on a bright reference star before slewing to the target star for the observation. This process is indicated with the solid blue line on the left side of Fig. 5.8. The right-hand side of the figure depicts the difference between a “set-and-forget”, or open-loop, observation strategy (dashed line) and a closed-loop observation strategy that uses continuous WFS&C (solid line).

In a perfectly stable setup, the coronagraph would be observing at c_0 for the duration of the observation, which is defined here with the total science exposure time t_{exp} . However, this is not what happens in the presence of disturbances on a real observatory structure due to thermal drifts from pointing changes, and dynamic loads such as thrusters or reaction wheels. These thermo-mechanical effects translate directly into dynamic wavefront errors. Between the beginning and the end of the science exposure time (from t_0 to t_1), the contrast therefore degrades and drifts by the contrast variation $\Delta c = c_1 - c_0$.

If the telescope is not stable enough on this timescale and the observatory is operating in open-loop, this contrast degradation Δc_{OL} will surpass the acceptable differential contrast for exoEarth imaging, defined as Δc_t .

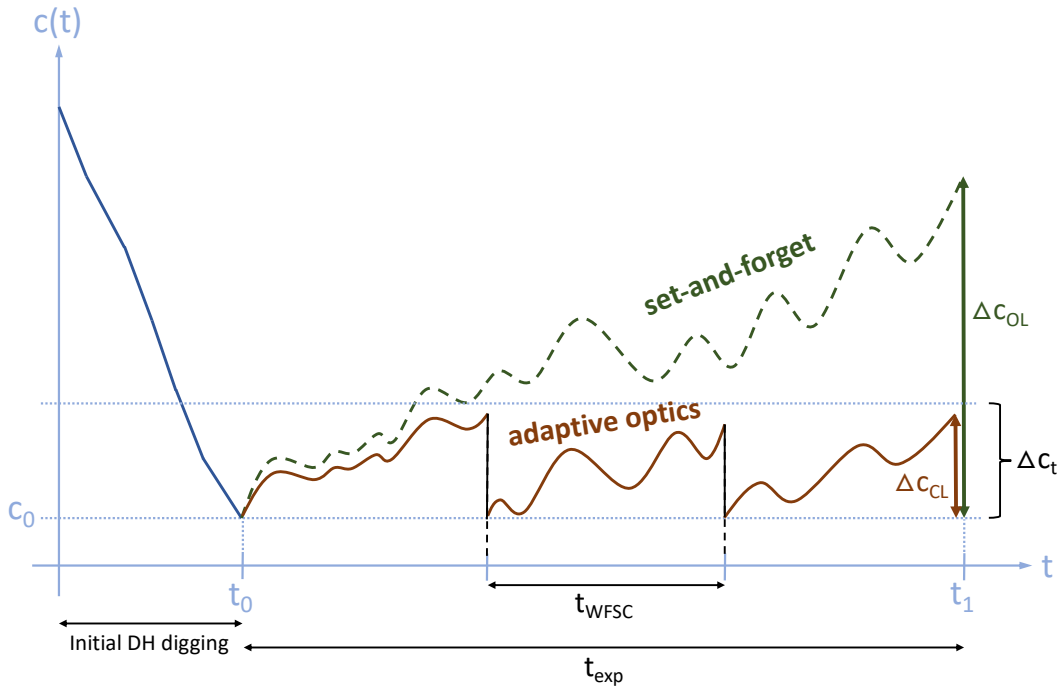


Figure 5.8: Set-and-forget WFS&C strategy (dashed line) with the associated open-loop contrast drift Δc_{OL} , compared to continuous WFS&C (solid line) with its associated close-loop contrast drift Δc_{CL} , as a function of time. Initially, a DH is established with DM commands computed from focal-plane WFS&C. After an instrumental raw contrast of c_0 is reached, the observation begins, with the goal to keep the contrast drift, or differential contrast Δc , within an acceptable range, Δc_t , for exoEarth detection.

The alternative approach of continuous WFS&C (i.e., similar to ground-based adaptive optics) allows us to recover the required contrast level regularly, on much shorter timescales than the total science integration time, which is indicated with the timescale associated with one WFS&C iteration, t_{WFSC} . By estimating and applying regular surface command updates for the DMs, the closed-loop differential contrast Δc_{CL} can be controlled such that it remains beneath the required threshold of Δc_t .

In the following sections, we show how to analytically express the WFE drift, d , in each of these observing scenarios, which can be expressed as the WFE variation $\Delta \epsilon$ over the respectively defining timescale t in each of the control strategies:

$$d = \frac{\Delta \epsilon}{t}. \quad (5.1)$$

In this chapter, we do not attempt exact derivations of the contrast and WFE values but instead we perform a dimensional analysis to provide general scaling laws for error budgeting and tolerancing purposes. While our focus is on the dynamic WFE requirements, we do *not* differentiate between different sources of WFE. The goal is to present an order-of-magnitude comparison between two fundamentally different wavefront control architectures (set-and-forget and AO).

Open-loop drift

When the DM settings from the DH solution are held statically during the full observation, all dynamic WFE that happen within the telescope and the HCI instrument need to be put in relation to the science exposure time t_{exp} , as it defines a natural limit for the effect and treatment of the occurring aberrations. The effect of WFE fluctuations on intermediate timescales on the same order as t_{exp} is integrated during the exposure time and therefore impacts the contrast in a direct fashion. Any drifts evolving on timescales larger than t_{exp} will be measured through focal-plane WFS and removed by the DMs in the same way as static errors in-between science exposures. Conversely, perturbations happening on timescales that are orders of magnitude faster than the scientific exposure time will cause a blurred halo in the coronagraphic images that could be removed in post-processing, in which case the remaining limitation is the associated

shot noise of the halo.

The goal is thus to determine the allowable WFE variation $\Delta\epsilon$ (corresponding to the contrast change Δc) over the time interval t_{exp} . This variation can then be converted into the tolerable OL drift d_{OL} per Eq. 5.1.

As we have shown in Chap. 2.2, the coronagraphic contrast is fundamentally a quadratic function of the wavefront (Eq. 2.9). This relationship was used in Chap. 4 to define the baseline contrast c_0 , which included uncorrected WFE from the segmented mirror, general testbed aberrations and the DM correction obtained with the DH algorithm. Our goal then was to focus on the validation of the PASTIS model, by introducing controlled aberrations into the testbed optics. Since the testbed is also impacted by natural drifts we cannot control, we needed to separate these contributions and therefore included all time-varying WFE in the term $c_0(t)$. In this chapter however, the goal is to establish requirements for the dynamic wavefront aberrations in general, irrespective of what causes them.

Consistent with Chap. 4, at the end of the DH algorithm process, we assume an observation that starts at an overall instrumental contrast $c = c_0$ at time $t = t_0$, which includes the effects from all physical aberrations like telescope phasing residuals and DM corrections. The contrast c_0 can therefore be expressed by dimensional analysis as being proportional to a quadratic wavefront error term ϵ_0 :

$$c_0(t_0) = \lambda^2 \epsilon_0^2. \quad (5.2)$$

This is a simplified, scalar version of the PASTIS equation given in Eq. 2.9, with a null static contrast floor. In this very simplified model, the observed real contrast c_0 is thus entirely described by the coronagraph sensitivity λ and a WFE ϵ_0 . Under this assumption, the only thing that can change the contrast is a variation in that WFE. As the wavefront drifts over the course of the observation, we can express the contrast at a later time $t = t_1$ as:

$$c_1(t_1) = \lambda^2 (\epsilon_0 + \Delta\epsilon)^2, \quad (5.3)$$

where $\Delta\epsilon$ is the WFE variation that induces the difference in contrast Δc between the beginning and the end of the science exposure time (see Fig. 5.8). This means we can write:

$$\Delta c = c_1 - c_0 \quad (5.4)$$

$$= \lambda^2 (\epsilon_0 + \Delta\epsilon)^2 - \lambda^2 \epsilon_0^2 \quad (5.5)$$

$$= 2\lambda^2 \epsilon_0 \Delta\epsilon + \lambda^2 (\Delta\epsilon)^2 \quad (5.6)$$

$$\simeq 2\lambda^2 \epsilon_0 \Delta\epsilon. \quad (5.7)$$

With the assumption that the minimum contrast is given by c_0 , the presence of additional residual aberrations ϵ_0 and $\Delta\epsilon$ yield a linear dependence of the contrast on $\Delta\epsilon$ as given by Eq. 5.6. In the small-aberration regime, with only small contrast fluctuations Δc as compared to c_0 , we can assume that $\Delta\epsilon$ is smaller than ϵ_0 . This allows us to approximate Δc with the linear expression given by Eq. 5.7.

Here, the variable Δc designates the same thing as in Chap. 4, except there we eliminated uncontrollable drifts to isolate the controllable aberrations, which we introduced on the segmented DM. By isolating the injected segment-level piston modes from the segmented DM in the contrast representations, we were able to study the impact of these specific modes. More generally, the goal in Chap. 4 was to validate the PASTIS model predictions independently of any time variations. Indeed, separate WFE aberration states were taken sequentially instead of relating them to specific timescales. We note that we could readily generalize the results obtained in Chap. 4 by scaling these wavefront variations to actual timescales. This is precisely the goal of this chapter: we focus on the tolerancing of a system evolving with time, but without introducing any additional aberrations. Therefore, the Δc here encompasses all the dynamic variations, for all aberrations, and the goal is to tolerance them using a dimensional analysis.

We choose to relate the instrumental contrast c_0 to the astrophysical planet-to-star flux ratio F by defining the ratio η :

$$\eta = \frac{c_0}{F}. \quad (5.8)$$

We solve Eq. 5.2 for $\epsilon_0 = \sqrt{c_0}/\lambda$, and substitute this together with Eq. 5.8 into Eq. 5.7 to obtain:

$$\Delta c = 2\lambda\sqrt{\eta F}\Delta\epsilon. \quad (5.9)$$

Now the goal is to specify a maximum tolerable Δc during the course of the science exposure time in the absence of a closed-loop adaptive optics correction.

In a purely idealistic static situation, the planet detection will be limited by the signal-to-noise ratio (SNR) due to photon noise from the speckle background (assumed uniform and calibrated out) at the level of contrast c_0 . Assuming that the planet would be detected at a SNR detection threshold κ_p (e.g., $\kappa_p = 10$ for a $10\text{-}\sigma$ level detection) we can define the maximum tolerable contrast variation Δc so that we are not dominated by the wavefront dynamical effects, but instead remain in a planet photon-limited scenario:

$$\Delta c = \frac{F}{\kappa_p}. \quad (5.10)$$

Inserting this into Eq. 5.9 and solving for $\Delta\epsilon$, we obtain a general expression for the WFE change:

$$\Delta\epsilon = \frac{\sqrt{F}}{2\sqrt{\eta}\lambda\kappa_p}. \quad (5.11)$$

With the WFE variation $\Delta\epsilon$ being defined, we now need an expression for the science exposure time t_{exp} to be able to determine the OL dynamical drift d_{OL} . Assuming that the sole noise source S in the data is the contrast-limited background photon noise N , we can do this with the definition of $\kappa_p = S/N$. The signal of the planet is given by the star's photon arrival rate, \dot{N}_s , times its flux ratio with the planet, F , integrated over the exposure time t_{exp} , written as $S = F\dot{N}_s t_{exp}$. The noise is proportional to the photon noise of the instrumental raw contrast $N = (c_0\dot{N}_s t_{exp})^{1/2}$, to a scaling factor. Since the derivations in this chapter are meant as general scaling laws for error budget and tolerancing purposes, we omit these scaling factors in the noise calculations. In practice, these scaling factors depend on the exact observing scenario (e.g., $\sqrt{2}$ for ADI reference subtraction (Marois et al. 2006), 1 for polarimetry (Kuhn et al. 2001), or $\sqrt{n_{ref}}$ for library-based subtractions (Soummer et al. 2012)). It is important to note that \dot{N}_s is an immediate function of the relative stellar magnitude of the observed star, m_V , and the telescope's collecting area as parametrized by its diameter D . We solve the resulting expression for κ_p , with substitution of Eq. 5.8, for the observation exposure time and obtain:

$$t_{exp} = \frac{\kappa_p^2 \eta}{\dot{N}_s F}. \quad (5.12)$$

The open-loop WFE drift requirement d_{OL} can then be calculated by limiting the WFE change over the time of a science exposure, and is thus given by $\Delta\epsilon/t_{exp}$, as discussed above:

$$d_{OL} = \frac{1}{\kappa_p^3} \frac{\dot{N}_s F^{3/2}}{2\eta^{3/2}\lambda}. \quad (5.13)$$

We can see that the only instrumental parameters in Eq. 5.13 are the choice of instrumental contrast parametrized by η , and the coronagraph sensitivity, given by λ . The less sensitive the coronagraph is to WFE (small values for λ), and the lower the ratio between the instrumental contrast and astrophysical flux ratio (small values for η), the easier it will be to stay within drift requirements. We note that choosing an instrumental contrast that is lower than the astrophysical brightness ratio makes $\eta < 1$, which contributes to a more relaxed WFE drift tolerance.

Closed-loop drift

The alternative approach to letting the telescope drift during the observation and counting on passive stability alone is to use continuous, or active, WFS&C throughout the observation sequence. A qualitative comparison between the two control strategies is shown in Fig. 5.8. Continuous WFS&C uses auxiliary wavefront sensors to measure misalignments, and then sends correction commands to the DMs in the coronagraphic instrument to correct for dynamic errors. In this way, the picometer-level stability does not have to be maintained over the total duration of the observation anymore, since the WFS&C loop keeps the contrast drift, the closed-loop residual, actively below target, $\Delta c \leq \Delta c_t$ (see Fig. 5.8). In this scenario, it is

sufficient to stabilize the system over the WFS&C time interval, t_{WFSC} (instead of the full science exposure time), which relaxes the WFE requirements now defined as $\Delta\epsilon/t_{WFSC}$. Here we assume a perfect controller with no lag, which means that the WFS&C time will be dominated by the time needed for sensing, t_{WFS} .

In the set-and-forget scenario, there is only one pertinent timescale for WFE variations, which is the science exposure time, t_{exp} . The WFS time t_{WFS} now adds an additional anchor point to the division between separate relevant timescales. We use this to derive the closed-loop drift requirement $d_{CL} = \Delta\epsilon/t_{WFS}$.

Assuming that the sensing error is driven by photon noise, the wavefront sensor error $\delta\epsilon_{WFS}$ is given by the WFE efficiency β over the square root of the flux on the WFS:

$$\delta\epsilon_{WFS} = \frac{\beta}{\sqrt{\dot{N}_s t_{WFS}}}. \quad (5.14)$$

Non-ideal wavefront sensors will generally have an efficiency $\beta > 1$, but using out-of-band sensing and/or predictive control can lead to values smaller than 1.

The maximum contrast variation that we can tolerate so that the planet detection remains limited by its own photon noise is given in Eq. 5.10. If we assume that this contrast variation is solely due to the wavefront sensor error $\delta\epsilon_{WFS}$ (perfect controller assumption), we can relate this contrast variation to the WFS error using dimensional analysis as $\Delta c = \lambda^2 \delta\epsilon_{WFS}^2$ and therefore we have:

$$\delta\epsilon_{WFS} = \frac{1}{\lambda} \sqrt{\frac{F}{\kappa_p}}. \quad (5.15)$$

Equating this expression with Eq. 5.14 allows us to solve for the wavefront sensing timescale t_{WFS} :

$$t_{WFS} = \frac{\beta^2 \lambda^2 \kappa_p}{\dot{N}_s F}. \quad (5.16)$$

Putting together the WFE change $\Delta\epsilon$ in Eq. 5.11 with the WFS timescale t_{WFS} in Eq. 5.16 yields the closed-loop drift requirement:

$$d_{CL} = \frac{1}{\kappa_p^2} \frac{\dot{N}_s F^{3/2}}{2\eta^{1/2} \beta^2 \lambda^3}. \quad (5.17)$$

Additionally to η and λ , the WFE drift is now also influenced by the WFS efficiency β . We observe how d_{CL} only scales with $\eta^{-1/2}$, which is the dependence on the ratio between instrumental contrast and astrophysical flux ratio of the planet, while a much greater influence is coming from β (squared) and the coronagraph sensitivity λ (cubed). This means that if the coronagraph is very robust, you can let the WFE drift for quite a while without degrading the contrast too much, especially with a good WFS efficiency.

Leveraging continuous WFS&C

The open-loop and closed-loop WFE drift expressions given in Eqs. 5.13 and 5.17 show a striking difference in their dependence on three core parameters of an HCI instrument. While d_{OL} is mostly determined by the raw contrast and the associated photon noise, d_{CL} relies to a higher degree on coronagraph robustness to misalignments. Fundamentally though, closed-loop operations significantly shorten the time span over which the WFE needs to be kept stable, with t_{WFSC} being orders of magnitude shorter than t_{exp} . We can now explore the advantages this brings in the temporal frequency domain.

Similar to the high-, mid- and low-spatial-frequency regime shown in Sec. 5.2, we can divide WFE timescales into three separate temporal frequency domains, set by t_{exp} and t_{WFSC} , as illustrated in Fig. 5.9. This description follows precursory work on error budgeting for JWST (Lightsey et al. 2014), and recent work on the same problematic for LUVOIR (Coyle et al. 2019b; Pueyo et al. 2019). We identify two factors that are responsible for WFE drift: thermal excitations leading to drifts as long as hours or minutes, and mechanical vibrations agitating telescope structures at frequencies of 10-100 Hz (Pueyo et al. 2019). The power of the ensuing temporal frequencies on which these WFE aberrations occur are characterized by their PSD, which has a qualitative example displayed in the top panel of Fig. 5.9. This red curve describes the optical distortions in the system without deploying any active control components like DMs. In the

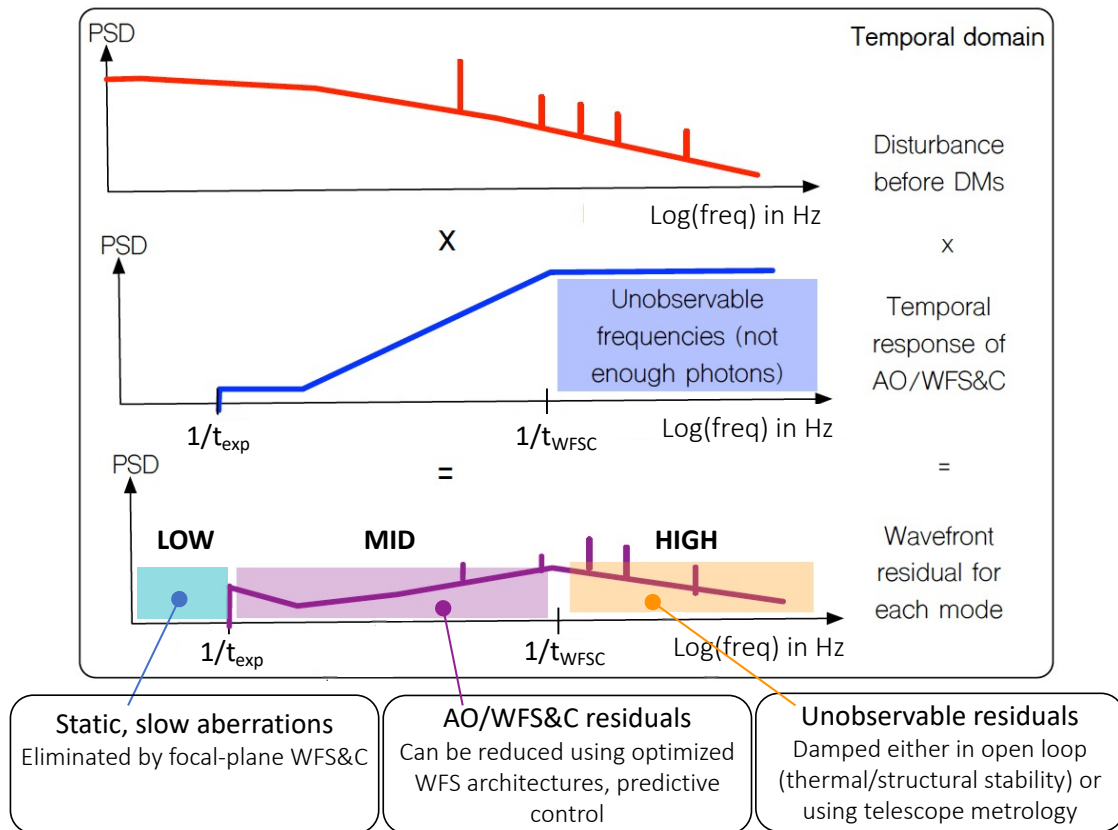


Figure 5.9: Temporal frequency domains as imposed by the two relevant timescales for observations with integrated AO: the science exposure time t_{exp} and the WFS&C time t_{WFSC} , which can be approximated by the time used exclusively for WFS, t_{WFS} , assuming a perfect controller. The PSD of the uncorrected WFE in the telescope is shown in the top panel, the filtering by the temporal response of the closed-loop WFS&C in the middle panel. The bottom panel indicates the resulting low-, mid- and high-temporal-frequency regimes that determine the residual WFE, and hence contrast performance, after correction. Details: see text. Adapted from Pueyo et al. (2019).

following, we want to relate the temporal frequencies of these aberrations to the temporal reference points defined by the frequencies $1/t_{exp}$ and $1/t_{WFSC}$.

In the temporal frequency domain, starting from the lowest frequencies, the first mark of concern is $1/t_{exp}$, which defines the point up to which any occurring WFE drifts are so slow that they can be considered static and thus do not contribute to the dynamic behavior of the contrast. Shortening the science exposure time shifts $1/t_{exp}$ toward the right on the middle panel in Fig. 5.9 with the goal to maximize this low-temporal-frequency domain while still allowing enough integration time for a sufficient SNR. Many coronagraphs are designed with a built-in rejection of certain aberration modes, which allows for longer integration times, and thus more signal, without a significant increase in WFE residuals. The more robust the coronagraph, the longer we can let the system drift without degrading our contrast too much.

For the set-and-forget scenario, $1/t_{exp}$ is the only relevant temporal frequency reference. Letting the structural steadiness of the telescope shoulder all the work to stabilize the observatory enough to keep the Δc within bounds would require extreme rigidity over very large spatial scales (Shaklan et al. 2006, 2005). This is why we are looking at the emergence of xAO technologies used in space in the future: they add a second temporal frequency up to which a range of drifts faster than $1/t_{exp}$ can be compensated with DMs, the maximal WFS&C frequency $1/t_{WFSC}$. The correction achieved with the AO system will not be perfect, but the continuously applied DM command updates informed by the wavefront sensors mean that the contrast drift can regularly be “reset” to an acceptable level smaller than Δc_t (Lyon & Clampin 2012), as illustrated by the solid brown line in Fig. 5.8. The limit $1/t_{WFSC}$ thus acts as the cutoff frequency of what is essentially a high-pass filter comprised by the WFS&C system. Finding an optimal cadence for WFS&C is a tradeoff between maximizing the SNR on the WFS by increasing t_{WFSC} and restricting the

WFE drift by limiting t_{WFSC} . Finding an ideal solution to this problem is contingent on the sensitivity of the wavefront sensor and coronagraphic robustness to WFE. Apart from these factors, faster cadence WFS&C is also favored by large telescope apertures, since their increased collecting area provides more favorable SNR on the WFS.

Any perturbations happening faster even than t_{WFSC} (which is bound to be faster than t_{exp}), with a frequency higher than $1/t_{WFSC}$, are unobservable residuals that remain uncorrected by the WFS&C loop. The aberrations in this high-temporal-frequency regime can be dampened with thermal or structural observatory stability, or with telescope metrology, but any net residuals will directly impact the science images. This leaves the mid temporal frequencies, falling right in between the observing and WFS&C timescales, as the main concern for active wavefront stabilization. The three relevant regimes in the temporal frequency domain are illustrated in the bottom panel in Fig. 5.9. Considering that WFS&C command updates can be generated on way faster intervals than the duration of science exposure times for the imaging of exoEarths at 10^{-10} contrast, a closed-loop observing strategy brings a significant easement in the requirements of WFE stability compared to telescope designs that envision to statically stabilize the full observatory structure over the course of an observation. In the following section, we perform a quantitative comparison of open-loop versus closed-loop requirements, in particular for segment-level aberration modes.

5.3.2 Drift requirements in a closed-loop WFS&C system

As introduced at the beginning of this chapter, numerical modeling as well as analytical treatment of the problem can be used to place quantitative WFE stability requirements on a HCI system aiming to reach 10^{-10} contrast. Conventional error budgeting approaches trickle this constraint down to all involved subsystems in an allocation that is *uniform in wavefront* across the board. Moreover, the general assumption is that the entirety of the observatory needs to stay sufficiently stable over the entire duration of the science exposure so as not to hamper the detection of a potential exoEarth.

A major result presented in Chap. 2 shows that an equal WFE allocation across all components is not optimal, and a better approach instead is to constrain the allocation to be *uniform in contrast*. This can be seen on the individual segment tolerancing presented in Chap. 2.4.2, applied to LUVUOIR A in Chap. 2.5.1 and displayed in Fig. 2.13. It shows that the narrow-angle APLC for example has a range of segment-level tolerances between 7 and 116 pm across the 120 segments, as set by the coronagraph sensitivity. While the segments on the inner rings still exhibit fairly constrained WFE limitations, the requirements on the outer segments seem significantly easier to meet under this consideration.

We noted before how the tolerancing analysis with PASTIS makes no assumption about the involved timescales, as the results it produces have a purely statistical meaning. The statistical per-segment requirements we obtain for example from the individual segment tolerancing implies that they will meet the target contrast as a *statistical mean* over many WFE states. The duration of time over which these states are realized is a concern of dynamical WFE tolerancing. As such, they can be interpreted as the static WFE requirements $\Delta\epsilon$ over the longest relevant timescale during an observation, which we saw in Sec. 5.3.1 is the science integration time t_{exp} . The inferred stability requirement of only a couple of pm over several hours is a truly intimidating goal, but it can be relaxed with the used of AO. In the same section, we learned how including a continuous WFS&C strategy not only reduces the timescales over which the telescope needs to be held passively stable, as shown in Fig. 5.8, but it also assigns more influence to coronagraph robustness for the purpose of relaxing these requirements.

We now aim to quantify the relaxation of WFE requirements as introduced by a continuous WFS&C strategy. To do that, we evaluate the expressions for open-loop and closed-loop drift requirements given by Eqs. 5.13 and 5.17. In both cases, we assume a star with an apparent magnitude of $m_V = 5$ and the spectral type A0V, its flux integrated over a 100 nm bandwidth from 500–600 nm, and the 15 m telescope architecture of LUVUOIR A with the narrow-angle APLC described in Chap. 2.5. The instrumental raw contrast is set to $c_0 = 10^{-10}$ and the differential target contrast to $\Delta c = 10^{-11}$, with an astrophysical brightness ratio for the planet of $F = 4 \times 10^{-11}$ and a planet SNR of $\kappa_p = 10$. For the closed-loop case, we assumed a theoretical WFS architecture with a sensing efficiency close to ideal, $\beta = 1$. We then calculate the segment-level WFE requirements both in OL and in CL, for an ensemble of six different localized Zernike




Figure 5.10: Gain associated with closed-loop WFS&C on a large, 15 m segmented telescope like LUVUOIR A, for a target contrast of $\Delta c = 10^{-11}$, separately for the inner segments on the left and outer segments on the right. The black, dashed horizontal line indicates a variation limit of 10 nm/min. Above this point, the analytical model for WFE drift might not be valid anymore, as the linear approximation made for the electric field relationship is suspended. The calculations use the WFE drift equations given in Eqs. 5.13 and 5.17. For details, see the main text.

modes on all 120 segments and their sensitivities as calculated in Sec. 5.2.2, and compare them.

The requirement relaxation associated with WFS&C is shown in Fig. 5.10, separately for the inner and outer segments of the LUVUOIR A telescope architecture. Operating in closed loop relaxes requirements for this observatory structures by at least a factor of ten, and very often by two orders of magnitude. It is remarkable that WFS&C techniques, which only involve minor modifications in an exoplanet imaging instruments (adding a WFS since DMs will already be present to reach the static contrast), can make it one hundred time easier to stabilize the large structure in a space telescope. In this analysis, the modes that gain the least (only a factor of four in requirement relaxation when compared to open-loop operations) are tip and tilt on the inner segments. In this example we thus have reduced the problem to a very small subset of pathological modes with relatively tight requirements at the telescope level. This can be mitigated by optimizing error budget allocations. Indeed, in practice one can give more room for these few modes to drift, as some less constrained modes will meet their allocations (which are larger than nanometers per minutes) with significant margins. However, this re-allocation between modes, that is somewhat architecture specific, is left to more detailed studies in the future. Because of the tremendous photon collecting ability of the 15 m primary mirror, even the tightest requirement would be of the order of 20 pm per minute, commensurate with expected segment-level thermal drifts for JWST (Lightsey et al. 2012). This highlights the tremendous potential associated with WFS&C during the course of long science exposures for future exoplanet imaging missions.

5.3.3 Perspectives for dynamic WFE tolerancing

With the above results, the advantage of using continuous wavefront correction becomes strikingly clear. However, the examples shown are based on an analysis of very generic aberration modes, and an order-of-magnitude model. They offer themselves very well for general studies, but we can use both more realistic derivations of the WFE drift, as well as a realistic modal basis to produce more practicable results. I was directly involved in the development of the above order-of-magnitude model during my PhD, and the resulting Eqs. 5.13 and 5.17 are sufficient to prove the advantage of closed-loop telescope controls over open-loop operations, as demonstrated with the results in Fig. 5.10. In this section however, we want to switch to a more advanced approach for the derivation of dynamic WFE drifts, by using a method developed by Pogorelyuk et al. (2021). We combine it with the thermo-mechanical tolerancing basis shown in Fig. 5.6, for which the static WFE tolerances in units of mK are shown in Fig. 5.7. With the work that follows here, we highlight the relevance of the PASTIS tolerancing model for a general statistical analysis of the influence of aberrations on the coronagraphic contrast, and how it can be combined with any temporal model to derive dynamic WFE drift requirements.

The segment requirement maps in Fig. 5.7 provide the equivalent of $\Delta\epsilon$ in units of allowable temperature change ΔT in mK, which we can use to express the drift requirement $d_{CL} = \Delta T/t_{WFS}$ in physical units. In order to constrain this drift quantitatively, we need to calculate an appropriate WFS timescale t_{WFS} . While the optimal scaling of the WFS time is out of scope of this thesis, there is an entire part of research that is looking at how to optimize the WFS&C parameters of a space-borne AO loop in order to limit residual WFE to an appropriate level. One method developed for this purpose is recursive estimation of the closed-loop contrast in a dynamic HCI instrument (Pogorelyuk et al. 2021).

This method supersedes the dimensional analysis approach in Sec. 5.3.1 as it provides scaling laws for the closed-loop variance based on an assumed open-loop WFE drift magnitude, stellar brightness, the WFE drift temporal PSD and detector noise. This allows us to calculate the optimal WFS time given a closed-loop target contrast. While CL operations significantly reduce the relevant timescale from t_{exp} to t_{WFS} , the resulting WFE drift requirement is also influenced by the different impact of instrumental parameters under CL conditions, as discussed in Sec. 5.3.1. This means that an optimization of the WFS time t_{WFS} needs to take into account that for example coronagraph robustness plays a stronger role in CL than in OL control strategies. Observing these differences between OL and CL (as is done by Pogorelyuk et al. (2021)), the transformation of OL WFE drift requirements to CL requirements is not exclusively a matter of dividing by a WFS time, but also by a scaling factor given by the optimization of all relevant instrumental parameters.

The main assumptions of the model developed by Pogorelyuk et al. (2021) are that of a discrete time sampling of t_{WFS} , that the open-loop drift follows Brownian motion which adds some variance at each iteration of the control loop, that we use a perfect controller and that we use an unbiased estimator. Under such conditions, the resulting WFE drift is balanced against the sensing noise to obtain an optimal WFS exposure time. We proceed by giving an example using a $m_V = 5$ star of spectral type A0V, with a flux integrated over a 100 nm bandwidth from 500–600 nm and on the 15 m-diameter aperture of LUVOIR A and for a target contrast of $\Delta c_t = 10^{-10}$. The results of such an iterative optimization of the photon information on the WFS as per Pogorelyuk et al. (2021, Sec. 2.3) are used to scale the static temperature gradient tolerances in Fig. 5.7 to dynamic drift requirements illustrated in Fig. 5.11. We assume a Zernike WFS that is not filtered by the FPM and sees the full telescope pupil (The LUVOIR Team 2019; Pueyo et al. 2019). In this case, the optimal WFS time using a batch-estimation is found to be $t_{WFS} = 4$ sec. We remind the reader that the drift requirements in temperature per time interval is given by the influence on contrast of the surface deformations in Fig. 5.6, which are caused by thermal gradients. The respective individual segment-level surface deformation maps are shown as figure insets in each global requirement map in Fig. 5.11, the latter given in units of mK per second.

The overall temperature stability requirements in Fig. 5.7 remain valid, but the allowable temperature *drifts* differ vastly between an open-loop and closed-loop case. For example, the temperature stability requirement of 215 mK for bulk segment heating translates to 1 mK/s for a 10 h open-loop observation sequence, but the same segments only have to be stabilized to 27 mK/s instead in closed-loop. Since the dynamic requirements result from a global scaling of the static results, this relaxation is observed for all segments and all modes on the primary telescope mirror.

While each of the local, thermal modes for which we calculate the dynamic tolerances that are shown in Fig. 5.11 ends up with different mK/s requirements, all of them take into account that the coronagraph used in this analysis (narrow-angle APLC) has a lesser sensitivity to segment-level aberrations on the outer segments. We have discussed this effect at length in Chap. 2.5 and Chap. 2.6. Since this is a product of the particular design of the coronagraphic masks, and the Lyot stop, we can see it in all segment-level tolerancing results for this particular coronagraph: from the piston-only tolerancing maps in Fig. 2.13, over static thermal requirement maps in Fig. 5.7, to the dynamic requirement maps of the same modal basis in physical units shown in Fig. 5.11. The segments along the edge of the telescope aperture are not constrained to the same deformation limits like the inner segments and the engineering effort that goes into keeping them stable thus does not need to be as strict. This comes to a big advantage of telescope metrology systems that rely on diagnostics of segment misalignments that are complementary to a low-order WFS (LOWFS) and OBWFS, like edge sensors. The segments on the outer rim of the telescope aperture do not have a full set of such sensors toward all sides, which will make them floppier by design,

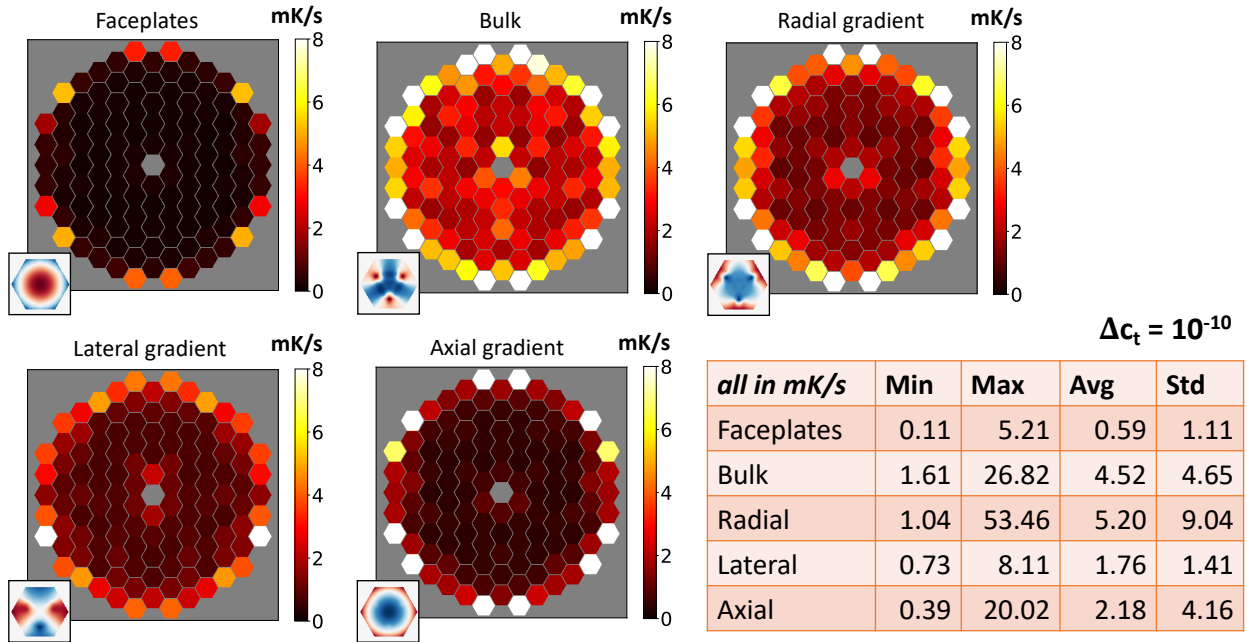


Figure 5.11: Dynamic temperature requirements for the narrow-angle APLC on LUVOIR A for a $m_V = 5$ star of spectral type A0V. The shown requirements give the acceptable temperature gradient in mK per second in order to keep the residual WFE under close-loop operations with WFS&C within a target contrast of $\Delta c_t = 10^{-10}$. These numbers assume a perfect controller and an unbiased estimator, and are calculated from the tolerance maps in Fig. 5.7 by using a scaling factor obtained with the method presented in Pogorelyuk et al. (2021). The inset figures below each requirement map indicate the used local mode from the surface plots shown in Fig. 5.6.

and having a lesser contrast sensitivity to such segments is thus beneficial.

The thermal aberration modes we addressed in this chapter represent thermal gradients that can be modeled as fairly slow drifts of several mK/s. An area of effort that needs to be addressed in the future is the determination of drift requirements for mechanical aberration modes, expressed in units of acceleration per time interval, or $\Delta g/t$. These manifest themselves on much faster timescales of 10–100 Hz as they originate from thruster and reaction wheel vibrations and therefore they need full dynamical timescale simulations. The ULTRA study team is tackling this problem in work efforts that carry on beyond my PhD with the implementation of the tools that I developed.

5.4 Conclusions

The PASTIS analytical tolerancing model is easily applicable to cases that go beyond the individual-segment tolerancing of piston errors presented in Chap. 2. The statistical implications of this model and its conclusions can be used for WFE tolerancing across a range of spatial frequencies, which covers global modes like Zernikes and sine wave ripples in the low- and high-spatial-frequency domain, as well as segment-level aberrations in the mid-spatial-frequency regime. Whichever modal basis is chosen, a PASTIS matrix can be generated to quantify the sensitivity of the coronagraph to WFE modes, used for analytical predictions of the mean contrast, of its variability and for WFE tolerancing studies. The development of the numerical optical simulators and the tolerancing package presented in Chap. 3 has enabled studies both within as well as outside of the direct scope of my PhD in order to quantify the allowable WFE when a particular coronagraphic contrast is to be achieved with a large telescope structure. An example presented in this chapter shows the calculation of WFE requirements for particular local segment modes obtained from thermo-mechanical modeling, which allow us to express allowable aberration levels in physical units.

The approach for coronagraph sensitivity analyses presented in this thesis makes no a priori assumptions about the strategy providing the necessary mechanical and thermal stability of the observatory. Its

derivations and conclusions are of purely statistical nature and need to be paired with an appropriate formulation of the temporal behavior of WFE to provide drift requirements that are commensurate with the start light suppression goals for exoEarth imaging. In this chapter we have confronted two functional approaches to enable the specification of telescope stability for this purpose. The first is the set-and-forget way in which the observatory needs near-perfect static behavior over the course of an observation to ensure low enough open-loop WFEs. The second one uses continuous WFS&C akin to extreme AO used routinely on the ground, where supplementary wavefront sensors and control components are used to regularly drive down the residual WFE.

We first described a dynamical drift tolerance formalism based on a dimensional analysis, which I was involved with as part of my PhD. This analytical, order-of-magnitude derivation for drift requirements in two WFS&C strategies (passive and active) shows a different dependence on key instrument parameters that influence the WFE drift in each scenario. The ratio of the instrumental contrast over the astrophysical flux ratio, η , defines the OL drift requirement with an exponent of $3/2$, while there is a first-order dependence on the coronagraphic sensitivity, λ . In CL operations however, it is the coronagraph sensitivity that influences the drift requirement as a cube, while the dependence on the contrast ratio is reduced to the order of only $1/2$. Moreover, the drift is influenced by the WFS efficiency β . Considering that coronagraphs for telescope designs like LUVOIR are designed with an inherent rejection to some aberration modes, the λ^{-3} and β^{-2} dependence in continuous WFS&C could allow us to exploit these instrumental characteristics for an improved stability level of aberrations in the future.

This simple formalism was more recently improved upon by that of Pogorelyuk et al. (2021) which we combined with the segment-level tolerancing results for thermo-mechanical modes obtained with PASTIS. By balancing the limitation on the WFS time t_{WFS} imposed by a limited photon flux and an ongoing WFE drift, this method computes a value for t_{WFS} that will allow us to meet a CL target contrast under consideration of all relevant instrumental parameters, like detector noise, WFS efficiency and coronagraph robustness. We were able to show that the tolerancing results obtained with PASTIS are fundamental to dynamic WFE tolerancing analyses, as they can be joined with methods optimizing the WFS&C loop to derive required WFE drift rates.

Taking all this into account, it follows that deploying a full AO system on a space-based telescope can provide significant relaxation of the WFE drift requirements by a factor of 10–100, as indicated by the results in Fig. 5.10. This is because the pertinent drift timescale in this control strategy is the time for a WFS&C iteration, rather than the full duration of a science exposure. This eliminates the reliance on a passive telescope structure that provides stability at the picometer-level without failure. While there is currently no space telescope that implements continuous WFS&C, RST is taking a step in that direction with “ground-in-the-loop” WFS&C, where the static contrast will be improved with focal-plane WFS supported by multiple DMs as well as “on-board” low-order fast WFS&C. Additionally, future space observatories can capitalize on an immense heritage of ground-based optical control experience from instruments like GPI and SPHERE, which is continuously evolving as we are preparing for new and more advanced imagers like SPHERE+, MagAO-X and KPIC. In the meantime, the individual subsystems that are necessary for closed-loop WFS&C in space are being continuously improved on. Wavefront sensing at the sub-nanometer level (Ruane et al. 2020; Steeves et al. 2020), robust coronagraph designs (Fogarty et al. 2020) and algorithms aimed at the continuous maintenance of WFE in particular (Redmond et al. 2021; Pogorelyuk & Kasdin 2019) have reached a component-level maturity commensurate with the needed specifications for a LUVOIR-like mission. On the road to fully develop these technologies, system-level demonstrations on laboratory testbeds and on-orbit will pave the way for these concepts to be implemented in space in the future.

The sensitivity analysis method developed during the course of this PhD is a cornerstone of these studies. It focuses on contrast allocations rather than top-down flows of uniform WFE attributions and provides a concrete way of defining stability requirements under consideration of the specific telescope geometry and coronagraph sensitivity. When paired with a suitable framework for the definition of wavefront control strategies on large, segmented telescopes, it enables a comprehensive evaluation of WFE requirements in a wide range of applications. It thus provides a fundamental basis for WFE tolerancing in future work.

6

CONCLUSIONS AND OUTLOOK

The challenges posed by the ambitious science goals of detecting exoEarths require on the one hand further understanding of the fundamental physics and image formation at high contrasts. In particular, some key questions remain open about the sensitivity to wavefront aberrations at the picometer level, polarisation effects, and the stability of control loops. On the other hand, we need extensive experimental validation of the methods and the models developed analytically. I dedicated my PhD work to both of these aspects.

The work I presented in Chap. 2 addresses the fundamental issue of segment-level wavefront error tolerancing. I have shown that we can in fact determine per-segment WFE tolerances in a purely analytical derivation and described them with a meaningful statistical model. Applying this analysis framework to the APLC designs on LUVOIR, I demonstrated that not all segments in the pupil have an equally tight constraint on their WFE levels. Not only does this let us derive concrete WFE requirements, but it provides us with a more palpable insight into the sensitivity of the coronagraph to specific aberrations. When we transform these sensitivities into an orthonormal modal basis, we can pinpoint specific modes that the contrast is more sensitive to than others. This led to some preliminary analysis results of the coronagraph sensitivity to aberrations induced by a continuous DM, rather than the segmented mirror the PASTIS formalism was originally developed for, and it harbors a range of interesting applications that should be explored in the future.

Among them, the more obvious is to utilize PASTIS for the tolerancing of Fourier-based aberrations on continuous surfaces, like those of a continuous face-sheet DM. Performing an analysis like this would allow us to specify the defining structure deformations for a monolithic telescope like RST, and verify the limitations deduced by traditional E2E modeling (Riggs et al. 2019; Nemati et al. 2017a). More than just pursuing a theoretical analysis, there is an interest in using the PASTIS eigenmodes, especially the ones on continuous DMs, in an optimized WFS&C scheme. There is a potential benefit to optimize the control loop for a contrast sensitivity criterion set by the coronagraph rather than purely reducing the Strehl ratio and the phase variance in the HCI system. In this way, we could make sure to enhance the results of the control loop specifically for the contrast performance. Additionally, one could identify and use a controller that is more efficient in suppressing the modes that degrade the contrast the fastest, or design a sensor that is by design more sensitive to these modes (Chambouleyron et al. 2021). Not only WFS&C could benefit from these results, but coronagraph design as well, as one can incorporate the sensitivities as found by PASTIS to create apodizers that increase the robustness to WFEs arising from segment misalignments (Leboulleux et al. 2021).

While PASTIS has been developed with the particular application to space-based telescopes in mind, it proposes itself naturally for the contrast sensitivity analysis of ground-based instruments, in particular segmented telescopes like the ELTs. The large optics on these observatories will suffer from global rigid-body deformations during slew times and mechanical adjustments. The tolerancing method treated in this thesis presents a tool that can be adjusted for analyses on such telescopes as well, provided it includes a model for long-exposure PSFs to account for the effects of the atmosphere. The representation of such long-exposure AO residuals has already been extensively considered in studies about HCI from the ground (Herscovici-Schiller 2018; Sauvage et al. 2010), and pairing them with the PASTIS model would provide a more complete tool for contrast-based tolerancing.

Studies like the ones mentioned above are enabled by the development of numerical tools like the

ones showcased in Chap. 3. In this chapter, I explain the work that went into building the software that enabled this entire thesis. This includes a modular, Python-based tolerancing package that performs the PASTIS tolerancing analysis, and a suite of numerical simulators for coronagraphic imaging with LUVOIR A, LUVOIR B, and HabEx. A particular focus was put on enabling the inclusion of externally-written optical simulators in the tolerancing tool, which I used to perform a sensitivity study on a NIRC*am* coronagraph on JWST. The results from this analysis show that the expected segment misalignments on the primary mirror of JWST are small enough to enable HCI at the 10^{-5} contrast level, the expected performance of these instruments (Perrin et al. 2018). The very same interface allowed me to perform simulations for a tolerancing analysis of the segmented DM on the HiCAT testbed. Much of my work included improving the testbed simulator for HiCAT, which I integrated with my WFE tolerancing tools in order to simulate cophasing requirements of the segmented DM on HiCAT. Due to the nature of the HiCAT simulator, which functions as an emulator of the real testbed infrastructure, performing these simulations meant to directly prepare the experimental validations I performed later.

The results of the experimental validations of the PASTIS tolerancing model on HiCAT are presented in Chap. 4. In fully remote operations of the Baltimore-based testbed while physically residing in Europe, I successfully measured an experimental contrast sensitivity matrix with respect to local piston aberrations on the segmented DM. After using this matrix to calculate the contrast eigenmodes and WFE tolerances for the individual segments, I confirmed through Monte Carlo simulations that these limits indeed lead to the envisioned contrast level. These experiments represent a validation of the underlying analytical model and its statistical meaning. I showed how to use this framework to compute WFE tolerancing limits on the segmented DM for various contrast levels, even if they lie beyond the current performance of HiCAT, and investigated how they change as a function of lateral LS misalignments. Acknowledging that PASTIS lets us make predictions about the contrast behaviour of an HCI instrument, given a set of known variances, it could find applications in combination with so-called DH maintenance algorithms (Redmond et al. 2021; Pogorelyuk & Kasdin 2019). These algorithms are not designed to create a high-contrast DH, instead they are deployed to keep the DH contrast from diverging too much over the course of an observation. DM dithers are used to increase phase diversity and estimate the speckle drift, and the commands used for sensing could be optimized with PASTIS by using poke patterns that the contrast is most sensitive to.

One of the problems that needed to be solved during the experimental validations is that the baseline contrast of a coronagraph does not remain stable over time. While I proposed a way to account for that during the experiments by isolating the contrast contribution from the segment aberrations alone, this leads to a larger question about sources of WFE in a large, segmented telescopes, in particular dynamic WFE sources. In Chap. 5, I first presented the results of a tolerancing analysis of very particular WFE sources coming from thermo-mechanically modeled segment-level aberration modes. This yields WFE requirements in terms of physical limits on the telescope, in this particular case how much of a temperature change is allowable if a certain target contrast is to be met. I further proceeded by putting these considerations in the context of applicable spatial and temporal scales that we expect for the evolution of dynamic WFE on a large telescope. Since PASTIS is a fundamental tolerancing tool with a purely statistical interpretation, it can in fact be used as a basis for sensitivity analyses across a wide range of spatial and temporal frequencies. Using the example of segment-level aberrations, I compared how the WFE requirement differs between the assumption of a telescope controlled in open-loop, or one with closed-loop WFS&C. The result is pretty clear: using a continuously running sensing and control loop on a large telescope relaxes the WFE requirements by at least a factor of 10, which justifies the complex systems currently envisioned for LUVOIR and HabEx, similar to extreme AO on the ground.

Indeed, space- and ground-based imaging research has long developed complementary aspects that are now mutually benefiting each other. Space-based instrument designs, not limited by atmospheric WFE, have advanced extremely precise coronagraph solutions to obtain the best performance with large apertures containing secondary struts or segmented primaries. With the progress of extreme adaptive optics, these complex apertures are now becoming the limitation of ground-based instruments and these coronagraphic techniques will be crucial to design the next generation of instruments for the ELTs. In turn, space-based instruments will soon be limited by very small WFE due to mirror quality or pointing errors. Decades of wavefront sensing and calibration research from the ground are now greatly informing

high-contrast space applications aboard future telescopes. With the RST launch predicted in the mid-2020s, we are anticipating the first space telescope to fly deformable mirrors, enabling active correction of the wavefront on-orbit. As it is a designated technology demonstration, the lessons learned from the operation of Roman CGI will be crucial to inform the concrete implementations of extreme AO in space, in the same way the segmented mirror technology of JWST is for LUVOIR. Such active sensing and control systems will be indispensable to enable the large, complex structures of segmented telescopes to achieve the WFE stability that is required for exoEarth imaging.

The work presented in this thesis contributes to these endeavors through the development of analytical models for statistical WFE tolerancing that relate the various influences of WFEs on a large, segmented telescope back to the coronagraphic contrast. In the end, what matters is the scientific output of our observations, and in the case of exoEarth imaging, an important metric for that is exoplanet candidate yield. Larger telescopes enable larger yield numbers, and to build such telescopes, we need to know how they will behave with respect to the performance of the HCI instrument. A tremendous effort has already been put in developing the individual subsystems that enable such complex observatories, and the ability to put these pieces together will in the end determine whether we can tackle the question of whether we are alone in the universe, or not.



COPHASING OF A SEGMENTED TELESCOPE TESTBED: JOST

After joining the Makidon lab team in late 2017, the main focus of my work was on laboratory experiments for the purpose of WFS&C on segmented telescopes. While the various activities on HiCAT were targeting in particular system-level operations of a segmented coronagraph, the JWST Optical Simulation Testbed (JOST) provided grounds to specialize on the WFS&C aspect. The synergy between the two testbeds in terms of goals and operations proved vital to start developing my expertise on hardware and software for high-contrast imaging laboratory experiments.

I inherited a fully aligned testbed from the previous lab team member Sylvain Egrou that simulated JWST with its most important degrees of freedom, the piston/tip/tilt controls of the segmented DM and misalignments of the secondary mirror, for which JOST uses a lens. Sylvain had previously measured wavefront errors in individual field points of the setup before the segmented DM installation, and demonstrated closed-loop linear control of the secondary lens. With the work presented in this chapter, adapted from a conference proceedings paper (Laginja et al. 2018) for SPIE Astronomical Telescopes & Instrumentation in Austin, Texas (USA) in 2018, I show the results of closed-loop WFS&C on the fully aligned JOST testbed, including an IrisAO DM. This was aided by my recalibration of the segmented DM flatmap in front of an interferometer, the installation of a faster camera and substantial simplifications and upgrades in the software control code.

Investigations like my work on JOST are relevant for WFS&C in the context of segmented telescopes. It allows us to develop and test new algorithms for the coarse and fine alignment of segments, which constitutes a vital part of a full HCI system.

James Webb Space Telescope Optical Simulation Testbed V: Wide-field phase retrieval assessment

A

I. Luginja, G. Brady, R. Soummer, S. Egron, C. Moriarty, C.-P. Lajoie, A. Bonnefois, V. Michau, E. Choquet, M. Ferrari, L. Leboulleux, O. Levecq, M. N'Diaye, M. D. Perrin, P. Petrone, L. Pueyo, A. Sivaramakrishnan

Appeared in
Proc. SPIE 10698, 106983N (2018)

ABSTRACT

The James Webb Space Telescope (JWST) Optical Simulation Testbed (JOST) is a hardware simulator for wavefront sensing and control designed to produce JWST-like images. A model of the JWST three mirror anastigmat is realized with three lenses in the form of a Cooke triplet, which provides JWST-like optical quality over a field equivalent to a NIRCcam module. An IrisAO hexagonally segmented mirror stands in for the JWST primary. This setup successfully produces images extremely similar to expected JWST in-flight point spread functions (PSFs), and NIRCcam images from cryotesting, in terms of the PSF morphology and sampling relative to the diffraction limit. The segmentation of the primary mirror into subapertures introduces complexity into wavefront sensing and control (WFS&C) of large space based telescopes like JWST. JOST provides a platform for independent analysis of WFS&C scenarios for both commissioning and maintenance activities on such observatories. We present an update of the current status of the testbed including both single field and wide-field alignment results. We assess the optical quality of JOST over a wide field of view to inform the future implementation of different wavefront sensing algorithms including the currently implemented Linearized Algorithm for Phase Diversity (LAPD). JOST complements other work at the Makidon Laboratory at the Space Telescope Science Institute, including the High-contrast imager for Complex Aperture Telescopes (HiCAT) testbed, that investigates coronagraphy for segmented aperture telescopes. Beyond JWST we intend to use JOST for WFS&C studies for future large segmented space telescopes such as LUVOIR.

A.1 Introduction

For large segmented telescopes, there is the need to actively control the telescope in order to achieve the optimum alignment and optical quality, bringing the primary from millimeter misalignments to fine alignments of nanometers. While controlled optics have become common on ground based telescopes, this technique will be extended to space for the James Webb Space Telescope (JWST). While its predecessor, the Hubble Space Telescope (HST), consisted of a mostly passive design with the exception of a variable defocus of the primary mirror, JWST will have 132 degrees of freedom between the primary and the secondary mirror. These will be initially aligned during the commissioning activities of the telescope, then maintained by periodic wavefront sensing and control (WFS&C) activities during its lifetime of at least five to ten years in order to maintain superb image quality (Acton et al. 2004, 2012; Knight et al. 2012b). The procedures for WFS&C on the JWST have been thoroughly tested in simulation and experiment (Barto et al. 2008) on the Testbed Telescope (TBT) at Ball Aerospace (Acton et al. 2006, 2007), a 1:6 scale model of the telescope that is equipped with the same degrees of freedom as the original, as well as the Integrated Telescope Model (ITM) software (Knight et al. 2012a). Most recently, WFS&C methods have been demonstrated on the integrated flight hardware (Acton et al. 2018; Lajoie et al. 2018). However, during the decade-plus development of JWST, WFS&C algorithms have continued to develop, and new advanced algorithms are worth investigating to expand the toolkit for alignment and maintenance of JWST.

The JWST Optical Simulation Testbed (JOST) at the Space Telescope Science Institute (STScI) provides a platform to test such algorithms for segmented mirror control, evaluating them for possible applications on JWST as well as on future space missions with active primary segmentation like the Large UV/Optical/IR Surveyor (LUVOIR) (Dalcanton et al. 2015; Bolcar et al. 2017). It is a simplified tabletop model of the JWST, as opposed to a high-fidelity scaled model like the TBT, but it is a close enough physical representation to model the key optical aspects. It is a supplement to existing verification and validation activities for independent cross-checks and novel experiments, not a part of the mission's critical path development process. In addition to exploring phase retrieval methods and implementing linear wavefront control over a wide field of view, JOST is used to develop staff expertise for commissioning and operations, conveniently being co-located at STScI with the Science & Operations Center (SOC) that will support commissioning and be responsible for operations of the JWST.

JOST is a three lens anastigmat, a refractive analogue to JWST's three mirror anastigmat. An aperture stop defines the system's pupil while the segmentation is provided by the planar segmented deformable mirror, whose segments can be controlled in piston, tip and tilt. The secondary lens (L2) that stands in as surrogate for JWST's secondary mirror is motorized in tip and tilt, and x, y and z translation. JOST has in total 59 motorized degrees of freedom, which are the most relevant ones for WFS&C maintenance activities. The setup design meets the requirement of an image quality of a minimum wavefront error of 40 nm rms at a wavelength of 638 nm over a field equivalent to one NIRCam module, and our latest measurements confirm that we meet this requirement, as we detect a minimum wavefront error rms of 15 nm.

This paper presents the optical characterization of JOST's full field of view after the successful fine alignment of L2 and DM, done previously. Our group presented a general overview of JOST in Perrin et al. (2014b). Its detailed optical design and several trade studies were presented in Choquet et al. (2014). The experimental implementation of the WFS&C on the testbed is described in Egron et al. (2016). The experimental results regarding the linear control of L2 are described in Egron et al. (2017a) and the alignment of the segmented deformable mirror is presented in Egron et al. (2017b). Before moving on to the implementation of WFS&C algorithms beyond the linearized algorithm for phase diversity (LAPD), we perform wide-field wavefront sensing on JOST. In section A.2 we give an overview of the testbed and its recent changes in hardware and software. In section A.3 we describe the wide-field wavefront sensing and present its results, and finally we summarize and conclude our findings in section A.4.




Figure A.1: JOST testbed layout. An off-axis parabola (OAP, not pictured) and a fold mirror put the laser beam launched from an optical fiber onto a steering mirror that is controlled in x and y by a stepper motor. This mirror illuminates the JWST-like pupil mask which directs the beam into the telescope simulator. The three lenses L1, L2 and L3 form a Cooke Triplet, a refractive analogue to the reflective three-mirror anastigmat (TMA) of the JWST. L2 acts as a surrogate for the JWST secondary mirror and is independently controllable by motors in tip and tilt, and x, y and z translation. Another fold mirror positions the beam on a subset of 18 segments of the IrisAO segmented deformable mirror, where all the segments can be independently controlled in piston, tip and tilt. A pupil imaging lens is attached to a flip-mount which allows for a fast change between pupil and focal plane imaging mode. The camera is mounted on a translation stage of 100 mm travel, which enables us to take phase-diverse data sets. The testbed can accommodate either a CMOS camera (shown here) for faster acquisition and smaller field of view, or a CCD camera providing the same field of view as a NIRCam module, with identical sampling.

A.2 Testbed description

An extensive description of the JOST optical design can be found in Perrin et al. (2014b) and Choquet et al. (2014). The updated current layout can be seen in Fig. A.1. The main components of JOST are a fiber launch, steering mirror for wide-field exploration, a JWST-like pupil mask, a telescope simulator made of three custom lenses, a segmented deformable mirror (DM) and a camera on a translation stage, to be able to provide focus-diverse images.

A.2.1 Key hardware components

The segmentation of the testbed, including gaps of the same size ratio like on JWST, is provided by an IrisAO segmented deformable mirror. The entire mirror has 37 independently controllable segments in piston, tip and tilt. A conjugated pupil mask with a hexagonal central obscuration and spiders defines the area of 18 segments that ultimately form the pupil. The DM can be controlled either by a GUI that is provided by the company, or by directly using the application programming interface (API) written in the C programming language, which we can call from within a Python wrapper. For best performance of the testbed, we need to have a baseline flat map of the DM, which is the configuration of the segments that gives the best optical quality of our data. The DM calibration provided by the vendor was for the DM oriented in the horizontal plane, facing upwards, while we use it standing upright in a mount facing horizontally, which makes all the segments sag forward and introduce large local tilts. A first flat map was created by using the GUI and checking the results directly with a Fizeau interferometer, in 2016. While this flat map showed a major improvement over the factory-defined flat position of the DM, we were able to

do an even more accurate calibration in late 2017 by adjusting each segment individually instead of using only global modes. The improvements led to an overall surface flatness of 10 nm rms over the entire pupil of 18 segments, creating a PSF with 39 nm rms wavefront error. The calibration maps and the PSFs we observe when the according flat map is put on the DM can be seen in Fig. A.2. This flat configuration can be improved further by running closed-loop iterations of LAPD on it, with the segment piston, tip and tilt being the controlled modes of this WFS&C experiment. Fig. A.3 shows how the overall wavefront error rms value drops from 40 nm to 16 nm after six iterations.




Figure A.2: Comparison of different flat maps for the segmented DM. *Top row:* Interferometer measurements of the DM surface with different flat map implementations in the IrisAO GUI and the resulting surface flatness rms. The red hexagon in the two left pictures denotes the JOST subaperture, while the third picture shows the JOST subaperture only. Grey hexagons indicate the position of dead segments in the engineering-grade device used on JOST. *Bottom row:* The corresponding PSFs to each flat map and their wavefront error rms taken on JOST with the IrisAO installed. *Left column:* Vendor-provided flat map where all values for piston, tip and tilt in the mirror GUI are zero. The DM was initially calibrated in the horizontal plane, lying down flat, while we use it standing upright in a mount, which makes all the segments sag forward and introduce a large local tilt. This leads to a PSF that shows many individual sub-PSFs, as the segments are not stacked properly. *Middle:* First custom flat map from 2016. The segments have been calibrated with masks that ignore the interfaces between the individual segments and only with global Zernike mode adjustments. The PSF is stacked, but still shows residual aberrations in the overall PSF structure. *Right:* Second and current custom flat map from late 2017. Instead of global modes, individual adjustments of piston, tip and tilt have been made on each segment. Further, the mask to define the JOST pupil also includes the discontinuities in the segment gaps, giving a better estimate of the final performance.

The telescope simulator consists of three custom made lenses that form a Cooke triplet, a refractive analogue to the JWST three mirror anastigmat, providing good optical quality over a wide field of view. During the alignment process of JOST, the third lens (L3) was added to the setup in reverse (Egron et al. 2016). The main effect of this is to reduce the size of the well-corrected field of view, but since it does not change the general alignment physics, and for convenience of use (fast readout, shorter exposure times, low noise) we have switched to a CMOS camera, whose smaller field of view matches the performance of the




Figure A.3: Closed-loop LAPD WFS&C experiment with the DM segmented controls as control modes. *Left:* RMS wavefront error vs. iteration number. *Right:* PSF at each iteration step. Iteration 1 is the initial configuration of the testbed, with each consecutive iteration of LAPD wavefront sensing and linear control of the DM providing a better overall quality of the system. The rms wavefront error drops to 16 nm after only six iterations.

current system well, we have decided to keep this configuration for now, until all wide-field infrastructure is finalized. Both cameras can be readily interchanged.

In order to be able to provide focus-diverse images which are needed for phase retrieval, the camera is mounted on a translation stage which provides a movement range of 100 mm in the z-direction (optical axis). The initial setup of JOST was using a CCD camera. However, the closed-loop experiments for wavefront sensing algorithms require 20 to 100 images be taken during each loop, and with the exposure times of one to ten seconds of the CCD camera the readout overheads would take a disproportionately long time compared to the wavefront sensing computations. We replaced the CCD with a ZWO monochrome CMOS camera that operates with exposure times between 0.5 and 500 ms, thus making the process of image acquisition considerably faster. The pixel size of the new camera is almost a third of the old size, 3.7 microns versus 9 microns, but a 2x2 pixel binning in the data reduction process leaves us with an effective pixel size of 7.4 microns, which means that our focused PSFs are still sampled with a factor slightly over 2, albeit not at the exact same sampling as NIRCcam, but this has no impact on algorithm development, and both camera setups remain available.

The before-mentioned piston, tip and tilt controls of the segmented DM over all 18 segments total 54 degrees of freedom on the DM. The secondary surrogate lens L2 can be remotely controlled with a stepper motor in tip and tilt, and x, y and z translation, yielding a total of 59 motorized degrees of freedom on JOST that are used in the linear control model (Egron et al. 2016). While JWST has a total of 132 degrees of freedom, the ones JOST misses are clocking, radius of curvature adjustment, and x and y translation of each individual segment, all of which are control modes whose major adjustments will happen during the initial six month commissioning process after launch. The degrees of freedom which are the same between JOST and JWST on the other hand are those with the largest optical influence functions during wavefront maintenance activities.

There is a steering mirror, movable in tip and tilt, positioned right before the JOST pupil mask. With this motorized mirror we have the possibility of directing the laser beam to off-axis positions on the detector and exploring field-dependent aberrations.

A.2.2 Software upgrades

Between October 2017 and April 2018, effort was put into restructuring the JOST software components. The first and crucial step was to migrate all prototype software to a version control system, to allow for a smoother and safer way of collaboration. Furthermore, all code involved in JOST data acquisition, data

reduction and hardware control have been translated to and extended in Python¹.

JOST is co-located with the High-contrast imager for Complex Aperture Telescopes (HiCAT; Soummer et al. 2018) and the two testbeds are taking advantage of each other's developments. The HiCAT experiment control software was recently rearchitected as a modern, clean Python package by a professional software engineer. Each instrument and hardware component has an interface that is easily accessible by simply installing the HiCAT package and importing it. Each hardware interface follows a simple object-oriented paradigm where the parent is an abstract class (e.g., "Camera"), which defines specific methods and implements a context manager. Context managers are important for hardware control because they will gracefully close the hardware even if the program crashes unexpectedly. The child classes implement the abstract methods such as `open()`, `close()`, `takeExposure()` with code for the specific camera. This keeps the scripts generic and means changing cameras will have little to no impact on the code. A thorough description of the HiCAT software infrastructure is given by Moriarty in these proceedings (Moriarty et al. 2018).

We installed and started using the HiCAT package on JOST. Since JOST uses the same type of hardware like HiCAT (same laser source and motor controllers, same camera type but different model), we only needed to update the configuration file of the package and were able to use the code as is. The modular structure of the code allows for very fast and clean generation of new scripts and implementation of new experiments. These changes are intended to push for best practices in astronomy coding, and incidentally move away from more traditional programming languages used in astronomy like IDL and Mathematica, providing a concise environment for the work done. This arrangement will facilitate JOST's role in providing a flexible multipurpose laboratory testbed for the testing and validation of independent phase retrieval techniques.

A.2.3 Previous wavefront sensing and control activities

In Egron et al. (2016), the authors describe the alignment of the three lenses of JOST with a phase diversity algorithm provided by the Office National d'Etudes et de Recherches Aérospatiales (ONERA) in France, and a linear optical control model. At that point, JOST did not yet include the segmented deformable mirror. Their results show a symmetric degradation of the wavefront when moving away from the optical axis. In the following paper by Egron et al. (2017b), a linearized algorithm for phase diversity (LAPD) (Mocœur et al. 2009) was used for the cophasing of the newly inserted segmented DM with the aligned testbed. In this algorithm, the pupil is made out of 18 hexagonal subapertures that simulate the effective JOST pupil consisting of the pupil mask and the DM segmentation, while the previous algorithm used for the lens alignment was working with a circular pupil without any obscuration. LAPD allowed for the alignment of the 18 mirror segments in piston, tip and tilt on each segment individually, and having both the mirror and the lenses aligned left the total wavefront error of JOST with an rms of under 40 nm.

Egron et al. (2017b) completed the full automation of the JOST testbed with regards to hardware control, data acquisition and reduction, wavefront sensing with an arbitrary phase retrieval algorithm and wavefront control with a linear coupling model. While the WFS&C in Egron et al. (2016) was implemented on a wide field of view, covering a range of $(-1^\circ, 1^\circ)$, the WFS&C after the addition of the segmented mirror was performed only on-axis.

A.3 Wide-field wavefront sensing with a segmented deformable mirror

A.3.1 Goals of wide-field WFS demonstration

The new goal is to expand the testbed capabilities to operation on a wide field for all degrees of freedom of JOST, which means we want to implement a wide-field approach to WFS&C on both the degrees of freedom of the DM (18 segments times 3 modes) and the motorized L2 variations (x, y and z translation plus tip and tilt). To achieve this we are seeking a validation of the L2 alignment with a hexagonally segmented pupil (as opposed to the round pupil in Egron et al. (2016)) and a closed loop WFS&C performance on an

¹<https://www.python.org>

extended field of view with the DM. In this paper, we present the results of wide-field wavefront sensing with LAPD and a characterization of the JOST field of view now that the DM is in place.

Since our camera changed from a CCD to a CMOS camera, our field has shrunk to $0.6^\circ \times 0.9^\circ$, but we extended our characterization beyond that by translating the camera radially away from the optical axis. With the motorized steering mirror in place it is easy to find the optical axis at any time and it enables us to iteratively explore a wide field of view.

A.3.2 Data acquisition and reduction

To be able to run any kind of phase retrieval, we need to make an in-focus image and at least one defocused image of the testbed configuration we are interested in. On JOST, the defocus diverse data is acquired by moving the camera on a translation stage; we move the camera by 93 mm to introduce a defocus of 4 waves (23.88 rad). The nominal field of view of the camera is $0.6^\circ \times 0.9^\circ$, so in order to get images out to 1.0° , we move the camera to different lateral offsets. A one-time shift is not enough though: the further away we move from the optical axis, the bigger the offset on the camera between the focused and defocused images will be, so the range of field points we can cover with one camera offset gets smaller with every step further outward. With five offsets, we cover a lateral distance from -0.3° to 1.02° and at one offset position, we move the steering mirror to 20 different distances from the optical axis and obtain 20 focused and defocused images each, as well as background images for both camera positions, thus probing the field in 100 different positions.

An automated Python script processes the images through the standard steps of stacking, background subtraction, centering, 2×2 binning (leaving us with 512×512 pixel images), bad pixel correction and normalization. The result are one focused and one defocused image that are consequently used by the wavefront sensing code to determine the wavefront aberrations at all points in the field of view. We obtain the overall wavefront error rms values through LAPD, which also creates wavefront maps which we then decompose with the Python package POPPY (Perrin et al. 2012) into individual Zernike modes from defocus (Z4; we follow the Noll convention (Noll 1976) for Zernike numbering) to primary spherical (Z11).

A.3.3 Wavefront sensing results

In a first step, we inspected the point-spread functions (PSFs) and wavefront maps obtained by the data acquisition and LAPD wavefront sensing. In order to clearly see an aberrated PSF, one has to go to a field point well beyond 0.5° . The rightmost PSF frame in Fig. A.4 shows the distorted PSF at 1.0° , and in this image, astigmatism is very clearly seen. The wavefront maps show the wide-field aberrations a bit earlier, for example at 0.5° . The center right wavefront image in Fig. A.4, at a distance of 0.5° , is starting to show a global tendency of the dark and bright wedges typical for astigmatism. This global wavefront patterns becomes more distinct when looking at the far right wavefront map in Fig. A.4, at a field point of 1.0° .

In Fig. A.5, we can see how the overall wavefront error changes as a function of distance from the optical axis of the testbed. The wavefront error is relatively uniformly scattered around 40 nm rms until a radial distance of 0.4° and it starts to rise significantly beyond 0.5° . This confirms that the JOST anastigmat has good optical quality in a field of view with a diameter of 1° . There are some discontinuities appearing in the data around 0.4° , 0.66° and 0.85° , which indicate the interface between two datasets that have been taken before and after a lateral detector shift as described in Sec. A.3.2.

Using the wavefront maps generated by LAPD, we decomposed each individual wavefront at each field point into the 11 first Zernike modes of the Noll convention. The three modes contributing the strongest to the overall wavefront error have been found to be Z4, Z6 and Z11 - defocus, 0° astigmatism and 3^{rd} order spherical aberration. They are shown as a function of field position in Fig. A.6, with an average rms error of 7 nm rms for defocus, 13 nm for astigmatism and 0 nm for spherical in the inner region of the field of view until 0.5° . Defocus reaches 45 nm rms at 1.0° , while the astigmatism reaches 20 nm and spherical reaches 18 at the same field distance. The results reflect the overall wavefront error distribution from Fig. A.5: the graphs are flat out to a distance of 0.4° , beyond which they continuously rise beyond 100 nm rms after 1.0° . The linear nature of the aberration modes confirm the linear dependence of the Zernike terms for a given field of observation as a function of the misalignment of L2, as it was demonstrated in Choquet




Figure A.4: PSFs (top) and wavefront maps (bottom) of four different points in lateral translation from the optical axis. The wavefront quality is very uniform out to a distance of 0.5° , the PSFs look very much alike. While the wavefront maps of the PSF at -0.2° and 0.0° look very similar, the map of the PSF at 0.5° is starting to show global aberrations, as inferred by the darker areas in the top right and bottom left parts of the pupil. The PSF at 1.0° away from the optical axis is very clearly aberrated, showing very strong astigmatism, which is also confirmed in the Zernike decomposition of the wavefront maps in Fig. A.6. The wavefront map at a field point of 1.0° distinctly shows astigmatism and defocus.

et al. (2014). While the big contributions of defocus and astigmatism repeat the results from Choquet et al. (2014) and Egron et al. (2016) (where the wavefront analysis was done without the segmented mirror in the system), the spherical aberration was not expected to be this dominant. It is not clear at this point what causes it, especially since we would expect to see more significant coma in the off-axis PSF positions instead.

A.3.4 Comparison to previous results without the segmented deformable mirror

In Choquet et al. (2014), figures 8, 9 and 10 show the design-predicted Zernike coefficients as a function of the field angle. Comparing our results to those, we can confirm the field-dependent rise of defocus and astigmatism. While Choquet et al. (2014) do show an increasing amount of spherical aberration in the wide field, it is not as dominant as the astigmatism while our results show the two to contribute equivalently to the wavefront error. In addition, we do not see coma showing up in our analysis, while we would expect to see some the further outwards we move.

Egron et al. (2016) have shown similar results in their figures 2 (experiment) and 4 (simulation), although the simulations have shown only significant defocus and astigmatism appearing in the off-axis wavefronts, and no other modes. In that experiment, there was no coma detected except in one of the corner PSFs, consistent with our non-detection in this work. The authors hypothesized that it was introduced by a flawed behavior of the L2 motors, since it showed up in only one of the four corners. This makes us confident that our detections of defocus and astigmatism in the present paper are real; however we are not able to tell at this point why the spherical aberration is so strong.

These findings support the further development of JOST into a multipurpose testbed that provides the possibility to implement different wavefront sensing and control techniques. With further work in the upcoming months, we will be able to provide new wide-field evaluations through the implementation of new WFS&C algorithms, like the JWST baseline Hybrid Diversity Algorithm (HDA), Geometric Phase Retrieval (GPR), Optimized Phase Retrieval Algorithm (OPERA) and Estimation of Large Amplitude Subaperture Tip-tilt by Image Correlation (ELASTIC).




Figure A.5: Total wavefront error rms of radially translated PSFs on JOST, as retrieved by LAPD, taking into account the shape and segmentation of the JOST pupil. The wavefront error is relatively uniformly scattered around 40 nm rms all the way out to 0.4° , with a significant rise after 0.5° , from where on it continuously grows to over 100 nm rms at 1.0° away from the optical axis. The discontinuities in the data at 0.4° and 0.85° are the limits between two data sets that are separated by a shift of the camera, as described in Sec. A.3.2. Since the optical properties of the testbed will be radially symmetric around the optical axis, this shows that JOST has very good optical quality in a wide field of $1^\circ \times 1^\circ$, centered on the optical axis of the system. Note that our baseline wavefront error rms here is 40 nm, while we demonstrated in Fig. A.3 that we can align the testbed to about 15 nm rms. While in Fig. A.3 we demonstrate that we have the ability to go down to 15 nm, we did not change the baseline alignment of the testbed to match our best alignment state of 15 nm rms in this current wide-field characterization.

A.4 Summary and Conclusions

JOST is a hardware simulator designed to test and validate wavefront sensing and control algorithms on segmented apertures like that of the JWST. In these proceedings, we presented the hardware and software updates performed on JOST since late 2017 and showed the characterization of the JOST wide-field with LAPD wavefront sensing.

One of the main hardware updates is the implementation of a new CMOS camera that allows for a faster image acquisition, but it reduced the overall field of view to about a quarter of the initial area. The second major hardware update is the new calibration flat map of the segmented deformable mirror, which was achieved by tweaking each individual segment in the pupil in piston, tip and tilt, until an overall surface error of 10 nm rms was achieved and integrated as the new baseline flat map on JOST. On the software side, we migrated JOST to common Python tools in the Makidon Optics Laboratory. We put all code on version control, translated control and analysis scripts from Mathematica and IDL to Python, and make use of the HiCAT Python package for hardware control. The updated testbed is now a modular setup for wavefront sensing and control experiments.

We presented wavefront sensing results with the currently implemented LAPD wavefront sensing algorithm, going one step further in the validation of the wide-field optics since the implementation of the DM. We showed the overall wavefront error on field points ranging from -0.3° to 1.0° and presented the contributions from different Zernike modes. The total wavefront error rms rises from a nominal 40 nm until a field point of 0.4° , where it starts to increase continuously, to over 100 nm at 1.0° . The three main Zernike




Figure A.6: Decomposition of the wavefront error of radially translated PSFs on JOST into individual Zernike modes. We probed the wavefront from piston (Z1) to spherical aberration (Z11) and are showing here the three most influential modes: defocus, 0° astigmatism and spherical aberration. The discontinuities in the data at 0.4° , 0.65° and 0.85° are the limits between two data sets that are separated by a shift of the camera, as described in Sec. A.3.2. As is the total wavefront error rms (see Fig. A.5), the Zernikes remain flat throughout the PSFs until 0.4° off the optical axis. Beyond that, their contribution to the total wavefront error grow continuously throughout all the data. From Choquet et al. (2014) and Egron et al. (2016) we would expect a lot of defocus, astigmatism and coma, but not as much spherical aberration. We are currently investigating where this is coming from.

contributors to the wavefront error are Z4, Z5 and Z11, which are defocus, astigmatism and spherical aberration. While the defocus and astigmatism were expected from the analysis in Choquet et al. (2014) and Egron et al. (2016), the spherical aberration has not occurred this strongly before and our upcoming work will investigate the cause of it. We expect further results of the JOST wide-field characterization by incorporating the Zemax interface within the JOST code to run simulations of the setup.

Now that the infrastructure is in place, with the L2 alignment and the IrisAO alignment done individually and a wide field evaluation and validation of the wavefront sensing, we can proceed to investigate wide-field control solutions and comparison of multiple phase retrieval techniques, namely the Hybrid Diversity Algorithm (HDA), Geometric Phase Retrieval (GPR), Optimized Phase Retrieval Algorithm (OPERA) and Estimation of Large Amplitude Subaperture Tip-tilt by Image Correlation (ELASTIC).

B

SYSTEM-LEVEL DEMONSTRATIONS OF HIGH CONTRAST ON A SEGMENTED APERTURE TESTBED: HiCAT

The experimental core work of my PhD were my contributions to the High-contrast imager for Complex Aperture Telescopes testbed, HiCAT. Instead of pursuing the demonstration of just a specific coronagraph or an individual WFS&C technique, HiCAT is exploring the interaction between the different components in an observatory to optimize the global system-level performance for exoplanet detection and characterization with segmented apertures. Through significant contributions to the hardware installations and alignment, software infrastructure and experimental processes, I saw the testbed reach an average dark-hole contrast on the 10^{-8} level, with an open-loop stability on the same order over the course of many weeks.

HiCAT contains a range of sensitive hardware components that need to be carefully controlled to use them in the sensing and control loops within experiments. Several deformable mirrors, off-axis parabolas and optical masks are aligned with in-built Michelson interferometers and many of them can undergo a fine-alignment through motorized stages. While there is no component on the testbed that I did not touch during my time in Baltimore, I focused in particular on the installation and calibration of the IrisAO segmented deformable mirror. Using a Fizeau interferometer, I established the optimal open-loop surface calibration over the 37 segments of this DM and wrote its control software. After successful installation on the HiCAT testbed itself, this work lead to the experimental validations of the contrast stability model I presented in Chap. 4.

Apart from preparing HiCAT for experiments to test and validate the tolerancing method I described at length in this thesis, I put in a lot of work into upgrading the testbed and defining procedures for its operations. While I was there, we essentially finished the testbed: we installed new mounts for the Lyot stop, FPM, apodizer and DM1, and deployed a whole range of different WFS&C techniques. I received the opportunity to present the results of our work on HiCAT in an oral presentation at the Spirit of Lyot conference in Tokyo, Japan, in October of 2019. The materials in this chapter aim to paint a fuller picture of the HiCAT science, operations and infrastructure I enabled through my active investment in the project.

B.1 History and purpose

The motivation for the HiCAT testbed is to provide an integrated solution for starlight suppression in combination with active WFS&C on a primary mirror geometry that contains segmentation, as well as a central obscuration and spiders (N'Diaye et al. 2013). Its base design aims to use an Apodized Pupil Lyot Coronagraph (APLC, see Chap. 1.3.1) and control with two continuous deformable mirrors with broadband light in ambient air. In the optical design, particular care was dedicated to Fresnel propagation effects in order to minimize amplitude-induced errors from out-of-pupil optics, in particular one of the two DMs. Such effects were required to not exceed a contrast level of 10^{-8} , one order of magnitude fainter than the nominal target contrast of HiCAT (10^{-7}). First light results of a bare assembly (no apodizer or DMs) were obtained in 2014 (N'Diaye et al. 2014), and shortly after the first coronagraphic images were taken with a classical Lyot coronagraph (CLC) and one single in-pupil DM (N'Diaye et al. 2015a). First results with WFS&C on the CLC were realized through the speckle nulling technique in 2016 and 2017 (Lebouilleux et al. 2016, 2017a).

A first dark-hole demonstration on the fully assembled system was performed in 2018 (Soummer et al. 2018), with an APLC installed together with both continuous DMs as well as the segmented DM, including a central obscuration and spiders. Control with a single DM through speckle nulling reached a monochromatic contrast of 1.7×10^{-6} , and 6.3×10^{-6} in 6% broadband light centered on 640 nm, in a half-circular DH. These results were the first to use a completely overhauled software infrastructure scripted almost entirely in Python, with continuous integration and autonomous testbed operations (Moriarty et al. 2018). This was the state of the testbed when I joined its efforts during my PhD. In the months and years that followed, I participated in obtaining WFS&C results on a circular dark hole like the results presented in Chap. 4.3.2 and Fig. 4.3.

There is an inherent reliance on the individual components in a testbed like HiCAT, for example a particular wavefront sensor, optical mask or controller. Each of them is crucial for the overall observatory system, but the purpose of a system-level demonstration is to establish a “modus operandi” when a set of such sub-components is connected into a system with the ultimate scientific goal to optimize exoplanet yield. On HiCAT, the proxy performance metrics for scientific output is the instrumental contrast, and its stability over various time scales, similar to what is outlined in Chap. 5.3. To this end, I have worked on optimizing the interaction between subsets of the system components, identifying limiting factors for the contrast and improving the overall testbed performance. This included writing parts of the optical simulator that is used in the HiCAT testbed emulator (see Chap. 3.3.1), as well as characterizing the contrast stability under the presence of environmental changes such as temperature and humidity fluctuations. Moreover, one of the features of HiCAT is the fast turnaround time for changes between different coronagraphic modes, which consist of different combinations of optical masks and components on the testbed. After briefly describing the overall testbed layout and the involved hardware in Sec. B.2, I go into more detail about said testbed modes and overall HiCAT operations in Sec. B.3.

B.2 Testbed optical layout and hardware

The HiCAT testbed is installed in a class 1000 cleanroom inside an ambient-air enclosure mounted on a floating optical table shown in Fig. B.1, left, and the open testbed is shown in Fig. B.1, right. The temperature and humidity inside the enclosure are controlled with a valve-regulated flow of dry air, injected through air diffusers to reduce turbulence. The current semi-unfolded layout of the testbed is displayed in Fig. B.2. The dotted-border regions in this schematic indicate the different observatory components as realized on HiCAT. The light source module on the top left is located in a cleanroom right next to the HiCAT room shown in Fig. B.1. It allows for a choice between a monochromatic laser diode at 640 nm and a supercontinuum laser source for broadband operations through a color filter wheel. Optical fibers connect the light source with the beam launcher installed on the testbed.

The first accessible pupil plane is populated with a mount for the pupil mask in which various laser-cut aperture mask designs can be mounted with a magnetic three-point mount. The following pupil plane after an unused focus contains the IrisAO 37-segment deformable mirror. In the next reimaged pupil plane,




Figure B.1: *Left:* HiCAT enclosure in its cleanroom, mounted on a floating optical table. The theodolite in the bottom left is used for fine alignments during component changes. *Right:* Frontal view of the HiCAT testbed with unmounted enclosure panels. The light source injection unit is located in the far right corner and is not visible in this image. The white wide cables belong to the three DMs installed on the testbed (the two Boston continuous DMs to the left and the IrisAO segmented DM to the right; DM1 and the segmented DM are seen from the back). The red cylindrical pieces are various CMOS camera models by ZWO that are used both as science and pupil cameras, as well as the detector for the Zernike wavefront sensor and as alignment cameras in the Michelson interferometers. The testbed mode in this image is a “segmented CLC” mode where the apodizer is replaced by a high-quality flat mirror.

there is the mount intended for the apodizer in the APLC. Due to the use of easily mountable bonding cells, the optic mounted in this plane can easily be exchanged for any new apodizer or a high-quality flat mirror for CLC operations. The next pupil plane accommodates the first of the two continuous-surface Boston Micromachines DMs (DM1) which is followed by the second such DM (DM2) mounted ~ 30 cm out of pupil. The refocused beam finds its way to the reflective focal-plane mask, where the light rejected by the coronagraph passing through the central hole travels into the low-order wavefront sensing system (LOWFS; Pourcelot et al. 2022, 2021, 2020). A flip-in mirror right before the FPM serves as a pickup for a phase retrieval camera (Brady et al. 2019, 2018). The light reflecting off the FPM and thus staying in the coronagraphic system passes through a transmissive Lyot stop in the next pupil plane before entering the equivalent of the science instrument. A beam-splitter divides the light between a pupil imaging arm and the focal-plane image on the imaging camera.

B.3 Operations and testbed modes

Over time, HiCAT has evolved into a versatile testbed that supports several different coronagraphic modes as well as WFS&C techniques. While the implemented sensing and control techniques allow us to do a quantitative comparison in terms of their performance, limiting factors, and advantages and disadvantages with regards to their implementation (Will 2021), the various coronagraphic installations, or “testbed modes”, are used to refine our understanding of the different limitation factors on a high-contrast testbed.

The basic operational mode of HiCAT today is to use the pair-wise estimator together with the stroke minimization controller to obtain dark holes with the process described in Chap. 4.3. This WFS&C strategy is used as the benchmark against which all other sensing and control methods are compared to. While the workhorse coronagraph of HiCAT is an APLC, including a real (controllable) segmented pupil, it is possible to perform all experiments on HiCAT in several different testbed modes, with different permutations of optical components installed. A visual representation of the current most important testbed modes is shown in Fig. B.3. By running WFS&C experiments in such different setups, we are exploring different parameters and performances in these modes. To be able to do such comparative work, we needed to define

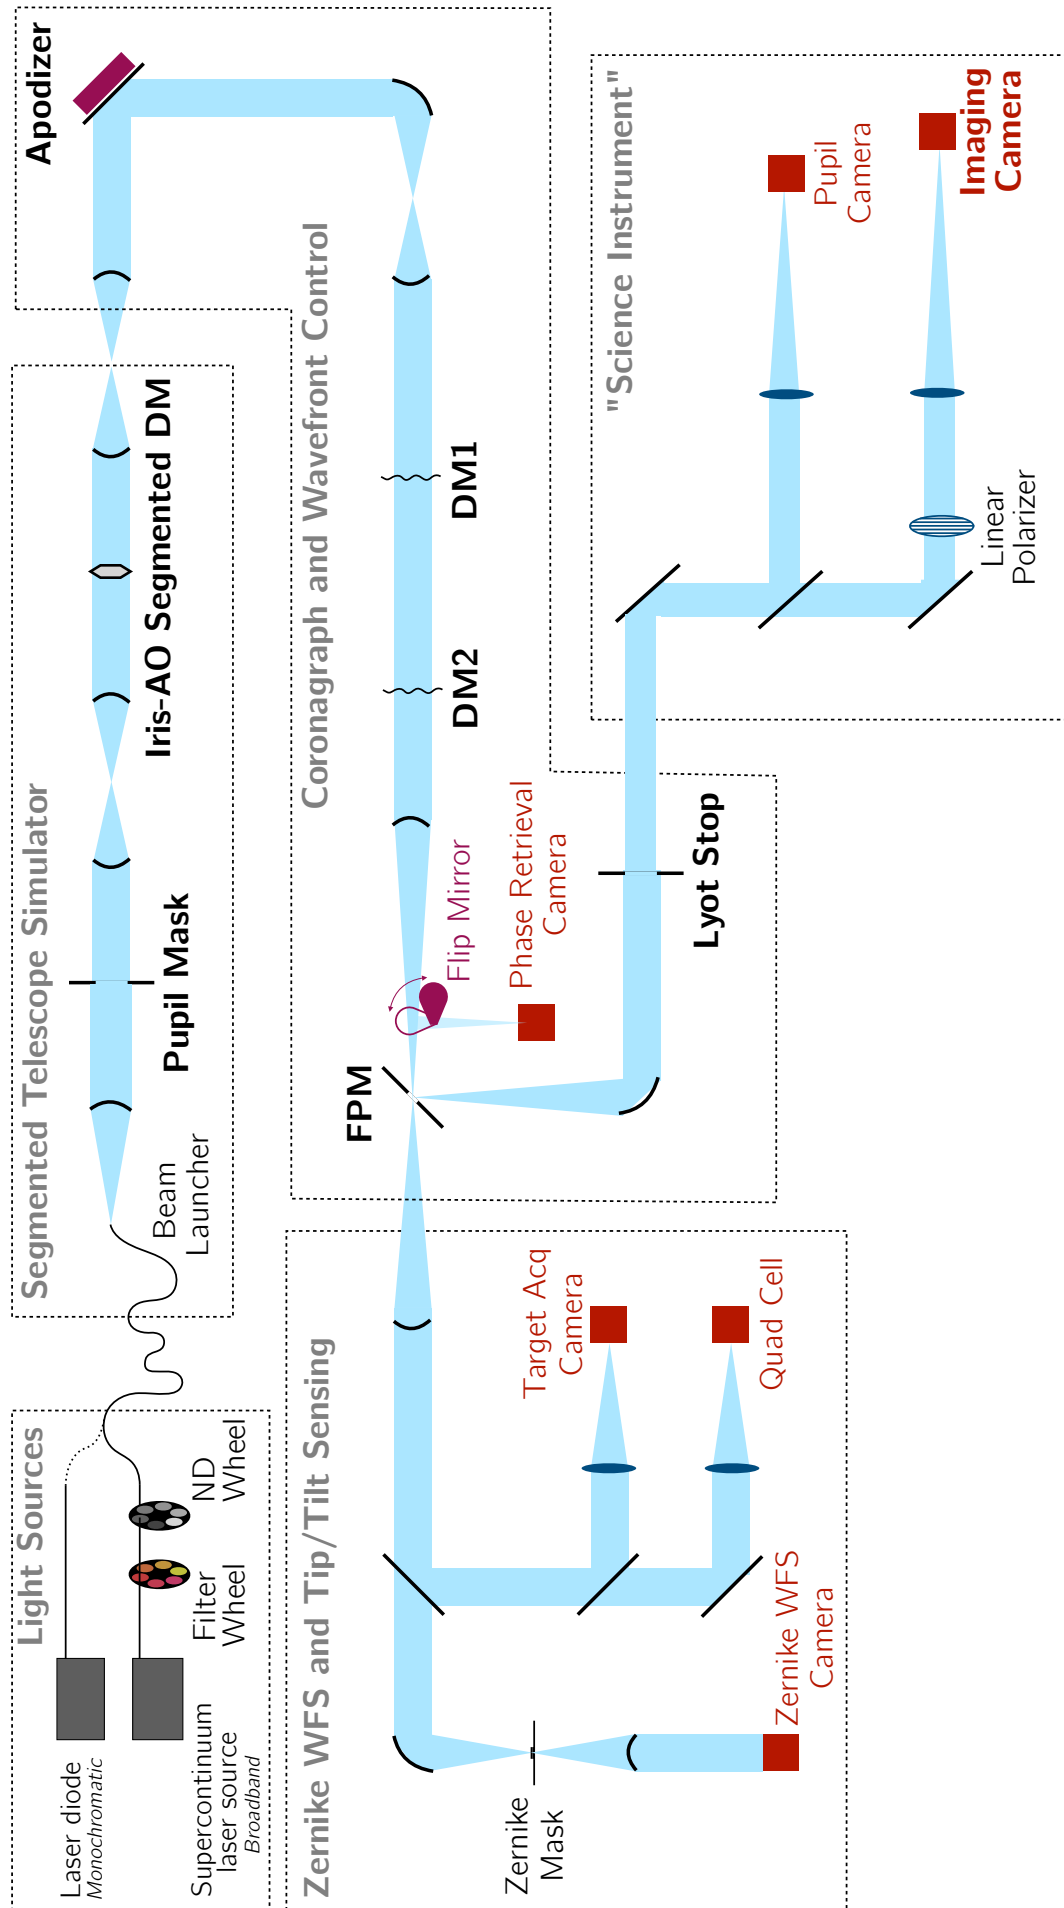


Figure B-2: HiCAT optical layout, unfolded into its relevant sub-systems. The light source is shown in the top left, followed by the segmented telescope simulator, and the coronagraph and WFS&C system. The light is split at the FPM with the light rejected by the coronagraph being injected into a LOWFS with fast tip-tilt control and a Zernike WFS. The imaging path provides pupil- and focal-plane images. For details, see main text. Credit: Marshall Perrin.




Figure B.3: Three of the different testbed modes in which HiCAT can be operated. The experimental DH results in the unsegmented modes (left and center) were obtained without the IrisAO segmented DM installed in the beam, the segmented mode (right) does include the segmented DM. All the DHs were obtained with a WFS&C loop using PW and stroke minimization. We note how the angular scale in the left two images is indicated in units of λ/D_{apod} , while the right DH is shown in units of λ/D_{LS} . At the time of writing, the best monochromatic contrast on HiCAT has been achieved in the segmented CLC mode (right), with an average contrast of 2×10^{-8} in a DH from $6 - 10\lambda/D_{LS}$. Not shown is the anticipated “full” testbed mode of HiCAT, the segmented APLC.

ways to switch back and forth between the different testbed modes without making each such mode swap take up too much time. I invested a lot of work in defining these procedures, including how to remove different hardware components, in which order to remove and to install them, and how to restore the performance once such a major hardware change is performed. The work I contributed to these efforts lead to a standard turnaround time between two testbed modes of no more than one day, including the basic testbed recalibration. With a finer recalibration performed one day after the mode swap, after the testbed has had time to settle, we recover the best testbed performance within 36 hours after starting the hardware change.

Over time, HiCAT has seen plenty of such mode swaps, with 11 of them alone happening in the time that I worked with the Makidon lab team. A timeline of the different HiCAT testbed modes and WFS&C activities from the beginning of 2017 until the end of 2022 is illustrated in Fig. B.4. The different testbed modes shown in Fig. B.3 (unsegmented CLC, segmented CLC, unsegmented APLC) are color-coded in this figure, with a distinction being made between *aplc1* and *aplc2*, because two different apodizer designs were used in these two modes. The small inset images are a visual representation of the pupil in each respective testbed mode. We note how the segmented CLC mode prior to 2019 was using a circular entrance pupil, while the segmented CLC mode after 2019 used a hexagonally outlined entrance pupil mask as described in Chap. 4.3.1.

Apart from the different hardware changes that happened over time, there were also a lot of changes and additions in the types of algorithms that are used to perform experiments on HiCAT. The first algorithm to dig a DH algorithm that was implemented and used until early 2019 was the speckle nulling technique (Lebouilleux et al. 2017a; Trauger et al. 2004). In 2018, the estimator called CORonagraphic Focal-plane waveFront Estimation for Exoplanet detection (COFFEE; Paul et al. 2014a, 2013) was successfully implemented on HiCAT using a segmented CLC with a circular pupil (Lebouilleux et al. 2020). In April of 2019, the now routinely used WFS&C strategy consisting of PW sensing and stroke minimization was implemented. With the end of 2020, the first experiments for the maintenance of a high-contrast DH have been enabled (Redmond et al. 2020, 2021, 2022). Early 2021 saw the addition of standard EFC (Give’on et al. 2007), Kalman Filters (Riggs et al. 2016; Groff & Kasdin 2013) and wavefront control through algorithmic differentiation (Will et al. 2022, 2021a,b).

The main interest of the operational developments of HiCAT and my involvement in them was to have a

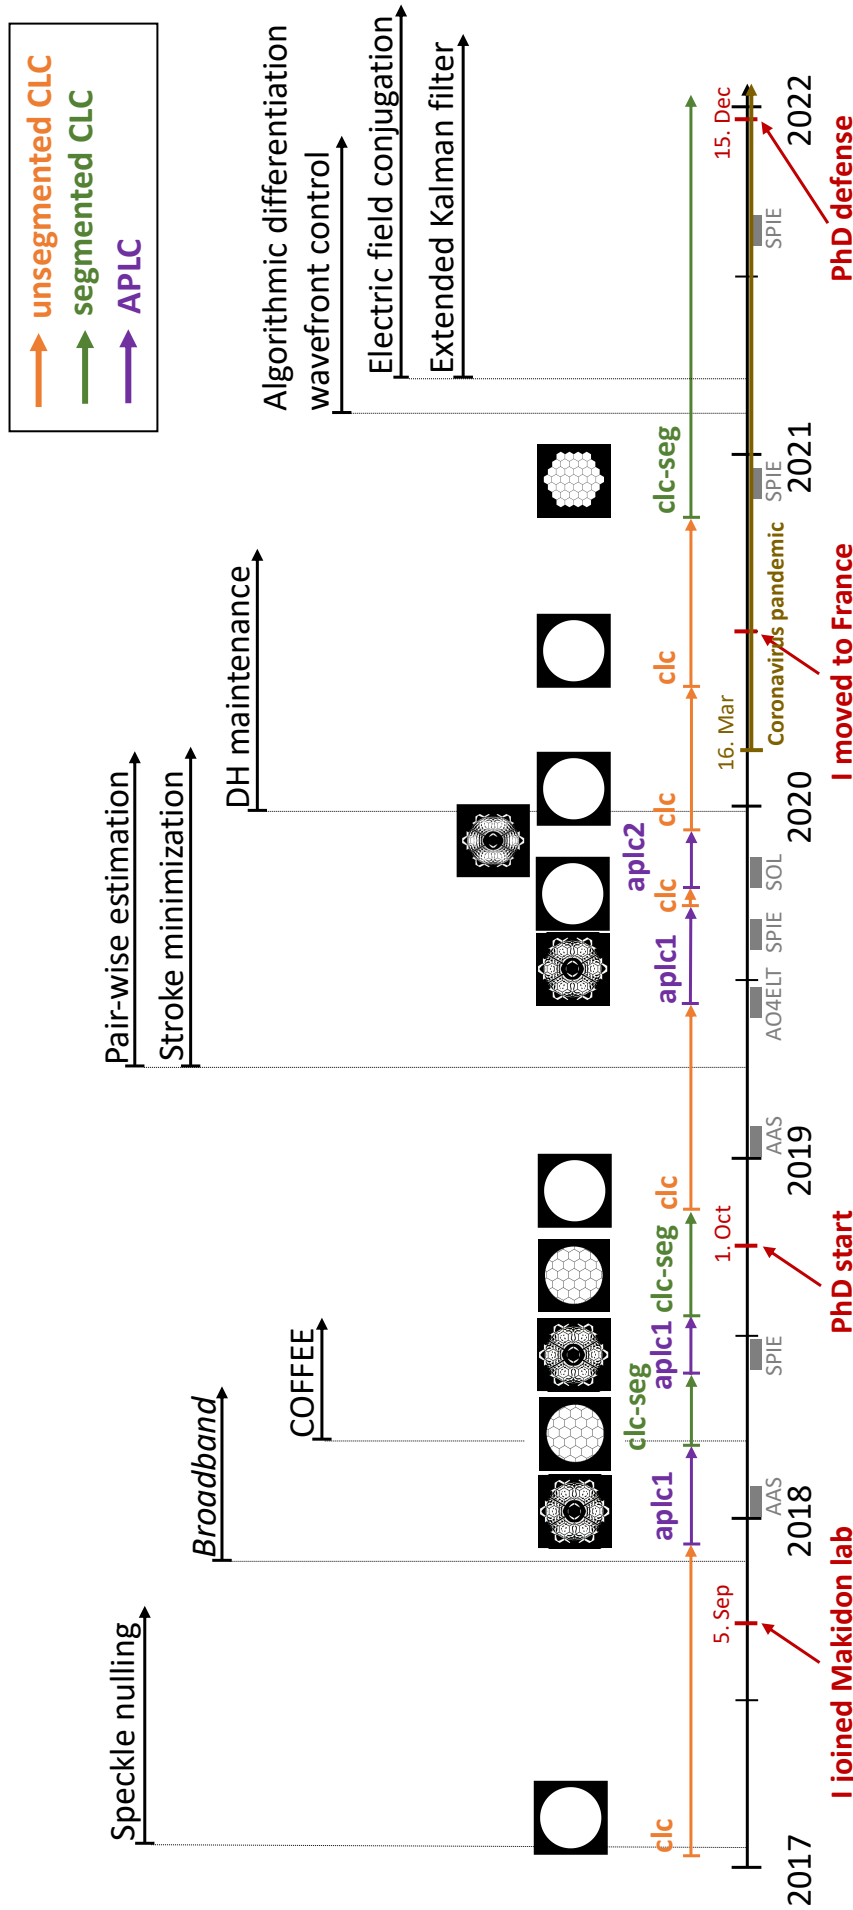


Figure B.4: Timeline of HiCAT operations from the beginning of 2017 until the end of 2022. The three tested modes shown in Fig. B.3 are color-coded and the small inset figures are a symbolic representation of these modes. We note that the aplc1 and aplc2 modes consisted of two distinct apodizer masks. The labels connected with vertical lines to the timeline indicated various DH algorithms and WFS&C techniques that were implemented on HiCAT in this time period. Overlaid in red are key dates of my PhD, together with the conferences I attended in grey. Also marked is the part of my PhD that I conducted during the Covid-19 pandemic. For details, see main text.

robustly functioning testbed on which we can compare results from one day to another, across many weeks and months. In my WFE tolerancing work that I presented in this thesis, a core point was to investigate whether the developed methods are in fact applicable to the real HiCAT system. I started conducting these experiments in late 2020, the successful results of which are presented in Chap. 4. By developing numerical simulations for the testbed, devising processes for hardware changes and installations, designing and running experiments on HiCAT, I have worked on all key aspects that constitute a high-contrast imaging testbed project.

BIBLIOGRAPHY

- Acton, D. S., Atcheson, P. D., Cermak, M., et al. (2004), *James Webb Space Telescope wavefront sensing and control algorithms*, Proc. SPIE, 5487, doi: 10 . 1117/12 . 551846
- Acton, D. S. & Knight, J. S. (2012), *Multi-field alignment of the James Webb Space Telescope*, Proc. SPIE, 8442, 84423C, doi: 10 . 1117/12 . 925004
- Acton, D. S., Knight, J. S., Chonis, T., et al. (2018), *Wavefront sensing and controls demo during the cryo-vac testing of JWST*, Proc. SPIE, 10698, 106983P, doi: 10 . 1117/12 . 2315396
- Acton, D. S., Knight, J. S., Contos, A., et al. (2012), *Wavefront sensing and controls for the James Webb Space Telescope*, Proc. SPIE, 8442, 84422H, doi: 10 . 1117/12 . 925015
- Acton, D. S., Towell, T., Schwenker, J., et al. (2007), *End-to-end commissioning demonstration of the James Webb Space Telescope*, Proc. SPIE, 6687, 668706, doi: 10 . 1117/12 . 731653
- Acton, D. S., Towell, T., Schwenker, J., et al. (2006), *Demonstration of the James Webb Space Telescope commissioning on the JWST testbed telescope*, Proc. SPIE, 6265, 62650R, doi: 10 . 1117/12 . 672133
- Ahn, K., Guyon, O., Lozi, J., et al. (2021), *SCEXAO: a testbed for developing high-contrast imaging technologies for ELTs*, Proc. SPIE, 11823, 1182303, doi: 10 . 1117/12 . 2594963
- Aime, C., Soummer, R., & Ferrari, A. (2002), *Total Coronagraphic Extinction of Rectangular Apertures Using Linear Prolate Apodizations*, A&A, 389, 334, doi: 10 . 1051/0004-6361:20020419
- Amara, A. & Quanz, S. P. (2012), *PYNPOINT: An Image Processing Package for Finding Exoplanets*, MNRAS, 427, 948, doi: 10 . 1111/j . 1365-2966 . 2012 . 21918 . x
- Bailey, V. P., Bottom, M., Cady, E., et al. (2018), *Lessons for WFIRST CGI from ground-based high-contrast systems*, Proc. SPIE, 10698, 106986P, doi: 10 . 1117/12 . 2313820
- Barto, A. A., Atkinson, C., Contreras, J., et al. (2008), *Optical performance verification of the James Webb Space Telescope*, Proc. SPIE, 7010, 70100P, doi: 10 . 1117/12 . 790483
- Baudoz, P., Boccaletti, A., Lacour, S., et al. (2014), *The high contrast imaging modes of MICADO*, Proc. SPIE, 9147, 91479E, doi: 10 . 1117/12 . 2057010
- Baudoz, P., Galicher, R., Patru, F., Dupuis, O., & Thijs, S. (2018a), *Status and Performance of the THD2 Bench in Multi-Deformable Mirror Configuration*, AO4ELT5, 1801. <https://arxiv.org/abs/arXiv:1801.06600>
- Baudoz, P., Galicher, R., Potier, A., et al. (2018b), *Optimization and performance of multi-deformable mirror correction on the THD2 bench*, Proc. SPIE, 10706, 107062O, doi: 10 . 1117/12 . 2314089
- Beaulieu, M., Martinez, P., Abe, L., et al. (2020), *High Contrast at Small Separation - II. Impact on the Dark Hole of a Realistic Optical Set-up with Two Deformable Mirrors*, MNRAS, 3926, 3914, doi: 10 . 1093/mnras/staa2106
- Belikov, R., Bryson, S., Sirbu, D., et al. (2018), *Design and performance analysis of a PIAACMC coronagraph on a segmented aperture*, Proc. SPIE, 10698, 106981H, doi: 10 . 1117/12 . 2314202

- Belikov, R., Pluzhnik, E., Witteborn, F. C., et al. (2012), *EXCEDE technology development I: first demonstrations of high contrast at $1.2 \lambda/D$ for an Explorer space telescope mission*, Proc. SPIE, 8442, 844209, doi: 10.1117/12.927218
- Beuzit, J.-L., Vigan, A., Mouillet, D., et al. (2019), *SPHERE: The Exoplanet Imager for the Very Large Telescope*, A&A, 631, A155, doi: 10.1051/0004-6361/201935251
- Bloemhof, E. E., Dekany, R. G., Troy, M., & Oppenheimer, B. R. (2001), *Behavior of Remnant Speckles in an Adaptively Corrected Imaging System*, ApJ, 558, L71, doi: 10.1086/323494
- Boccaletti, A., Chauvin, G., Mouillet, D., et al. (2020a), *SPHERE+: Imaging Young Jupiters down to the Snowline*, White paper submitted to ESO. <https://arxiv.org/abs/arXiv:2003.05714>
- Boccaletti, A., Folco, E. D., Pantin, E., et al. (2020b), *Possible Evidence of Ongoing Planet Formation in AB Aurigae - A Showcase of the SPHERE/ALMA Synergy*, A&A, 637, L5, doi: 10.1051/0004-6361/202038008
- Boccaletti, A., Lagage, P.-O., Baudoz, P., et al. (2015a), *The Mid-Infrared Instrument for the James Webb Space Telescope, V: Predicted Performance of the MIRI Coronagraphs*, PASP, 127, 633, doi: 10.1086/682256
- Boccaletti, A., Thalmann, C., Lagrange, A.-M., et al. (2015b), *Fast-Moving Features in the Debris Disk around AU Microscopii*, Nature, 526, 230, doi: 10.1038/nature15705
- van der Boekel, S., Colbert, S. C., & Varoquaux, G. (2011), *The NumPy Array: A Structure for Efficient Numerical Computation*, Computing in Science Engineering, 13, 22, doi: 10.1109/MCSE.2011.37
- Bohn, A. J., Kenworthy, M. A., Ginski, C., et al. (2020), *Two Directly Imaged, Wide-Orbit Giant Planets around the Young, Solar Analog TYC 8998-760-1*, ApJ, 898, L16, doi: 10.3847/2041-8213/aba27e
- Bolcar, M. R. (2019), *The Large UV/Optical/Infrared (LUVOIR) surveyor: engineering design and technology overview*, Proc. SPIE, 11115, 111150P, doi: 10.1117/12.2530456
- Bolcar, M. R., Aloezos, S., Bly, V. T., et al. (2017), *The Large UV/Optical/Infrared Surveyor (LUVOIR): Decadal Mission concept design update*, Proc. SPIE, 10398, 1039809, doi: 10.1117/12.2273848
- Borde, P. J. & Traub, W. A. (2006), *High-Contrast Imaging from Space: Speckle Nulling in a Low-Aberration Regime*, ApJ, 638, 488, doi: 10.1086/498669
- Boss, A. P. (1997), *Giant Planet Formation by Gravitational Instability*, Science, 276, 1836, doi: 10.1126/science.276.5320.1836
- Bowman, L. M., Belvin, W. K., Komendera, E. E., Dorsey, J. T., & Doggett, B. R. (2018), *In-space assembly application and technology for NASA's future science observatory and platform missions*, Proc. SPIE, 10698, 1069826, doi: 10.1117/12.2311413
- Bradley, L., Sipőcz, B., Robitaille, T., et al. (2020), *astropy/photutils: 1.0.1*, Zenodo, doi: 10.5281/zenodo.4049061
- Brady, G. R., III, P. P., Luginja, I., et al. (2019), *Phase-retrieval-based wavefront metrology for high-contrast coronagraphy: 2. Reconstructions through a shaped pupil apodizer*, Proc. SPIE, 11117, doi: 10.1117/12.2530482
- Brady, G. R., Moriarty, C., Petrone, P., et al. (2018), *Phase-retrieval-based wavefront metrology for high contrast coronagraphy*, Proc. SPIE, 10698, doi: 10.1117/12.2314368
- Brandl, B., Bettonvil, F., van Boekel, R., et al. (2021), *METIS: The Mid-Infrared ELT Imager and Spectrograph*, The ESO Messenger vol. 182, pp. 22-26, 5 pages, doi: 10.18727/0722-6691/5218

- Cady, E., Balasubramanian, K., Gersh-Range, J., et al. (2017), *Shaped pupil coronagraphy for WFIRST: high-contrast broadband testbed demonstration*, Proc. SPIE, 10400, 104000E, doi: 10.1117/12.2272834
- Cantalloube, F., Mouillet, D., Mugnier, L. M., et al. (2015), *Direct Exoplanet Detection and Characterization Using the ANDROMEDA Method: Performance on VLT/NaCo Data*, A&A, 582, A89, doi: 10.1051/0004-6361/201425571
- Carbillet, M., Bendjoya, P., Abe, L., et al. (2011), *Apodized Lyot Coronagraph for SPHERE/VLT: I. Detailed Numerical Study*, Exp. Astron., 30, 39, doi: 10.1007/s10686-011-9219-4
- Carlomagno, B., Delacroix, C., Absil, O., et al. (2020), *METIS High-Contrast Imaging: Design and Expected Performance*, JATIS, 6, 035005, doi: 10.1117/1.JATIS.6.3.035005
- Carlotti, A., Vanderbei, R., & Kasdin, N. J. (2011), *Optimal Pupil Apodizations of Arbitrary Apertures for High-Contrast Imaging*, Opt. Express, 19, 26796, doi: 10.1364/OE.19.026796
- Caswell, T. A., Droettboom, M., Lee, A., et al. (2020), *matplotlib/matplotlib: REL: v3.3.3*, Zenodo, doi: 10.5281/zenodo.4268928
- Chambouleyron, V., Fauvarque, O., Sauvage, J.-F., et al. (2021), *Variation on a Zernike Wavefront Sensor Theme: Optimal Use of Photons*, A&A, 650, L8, doi: 10.1051/0004-6361/202140870
- Chanan, G., Ohara, C., & Troy, M. (2000), *Phasing the Mirror Segments of the Keck Telescopes II: The Narrow-Band Phasing Algorithm*, Appl. Opt., 39, 4706, doi: 10.1364/AO.39.004706
- Charbonneau, D., Brown, T. M., Latham, D. W., & Mayor, M. (2000), *Detection of Planetary Transits Across a Sun-like Star*, ApJL, 529, L45, doi: 10.1086/312457
- Chauvin, G., Lagrange, A.-M., Dumas, C., et al. (2004), *A Giant Planet Candidate near a Young Brown Dwarf. Direct VLT/NACO Observations Using IR Wavefront Sensing*, A&A, 425, L29, doi: 10.1051/0004-6361:200400056
- Chilcote, J., Konopacky, Q., de Rosa, R. J., et al. (2020), *GPI 2.0: upgrading the Gemini Planet Imager*, Proc. SPIE, 11447, 114471S, doi: 10.1117/12.2562578
- Choquet, É., Levecq, O., N'Diaye, M., Perrin, M. D., & Soummer, R. (2014), *James Webb Space Telescope Optical Simulation Testbed II: design of a three-lens anastigmat telescope simulator*, Proc. SPIE, 9143, 91433T, doi: 10.1117/12.2056931
- Chun, M., Toomey, D., Wahhaj, Z., et al. (2008), *Performance of the near-infrared coronagraphic imager on Gemini-South*, Proc. SPIE, 7015, 70151V, doi: 10.1117/12.787683
- Ciliegi, P., Agapito, G., Aliverti, M., et al. (2021), *MAORY: A Multi-Conjugate Adaptive Optics Relay for ELT*, The ESO Messenger vol. 182, pp. 13-16, 4 pages, doi: 10.18727/0722-6691/5216
- Claudi, R. U., Turatto, M., Gratton, R. G., et al. (2008), *SPHERE IFS: the spectro differential imager of the VLT for exoplanets search*, Proc. SPIE, 7014, 70143E, doi: 10.1117/12.788366
- Codona, J. L. (2012), *Theory and application of differential OTF (dOTF) wavefront sensing*, Proc. SPIE, 8447, 84476P, doi: 10.1117/12.926963
- Codona, J. L. & Doble, N. (2012), *Experimental evaluation of differential OTF (dOTF) wavefront sensing*, Proc. SPIE, 8447, 84476R, doi: 10.1117/12.927025
- Codona, J. L., Kenworthy, M. A., Hinz, P. M., Angel, J. R. P., & Woolf, N. J. (2006), *A high-contrast coronagraph for the MMT using phase apodization: design and observations at 5 microns and $2\lambda/D$ radius*, Proc. SPIE, 6269, 62691N, doi: 10.1117/12.672727

- Cornelissen, S. A., Hartzell, A. L., Stewart, J. B., Bifano, T. G., & Bierden, P. A. (2010), *MEMS deformable mirrors for astronomical adaptive optics*, Proc. SPIE, 7736, 77362D, doi: 10.1117/12.857296
- Coyle, L., Knight, J. S., Hicks, B., et al. (2020), *Progress towards hardware demonstrations of critical component-level technologies for ultra-stable optical systems*, Proc. SPIE, 11443, 1144323, doi: 10.1117/12.2560185
- Coyle, L., Knight, S., Barto, A., et al. (2019a), *Ultra-Stable Telescope Research and Analysis (ULTRA)*, Tech. rep., NASA Roses, https://nspires.nasaprs.com/external/viewrepositorydocument/cmdocumentid=679697/solicitationId=%7BE879B8CD-FC43-BAB4-679E-39B4C90D567E%7D/viewSolicitationDocument=1/17SLSTD170003FRBall1201904Public_pdf
- Coyle, L. E., Knight, J. S., & Adkins, M. (2018), *Edge sensor concept for segment stabilization*, Proc. SPIE, 10698, 1069869, doi: 10.1117/12.2312224
- Coyle, L. E., Knight, J. S., Pueyo, L., et al. (2019b), *Large ultra-stable telescope system study*, Proc. SPIE, 11115, 111150R, doi: 10.1117/12.2525396
- Crossfield, I. J. & Troy, M. (2007), *Segment aberration effects on contrast*, Appl. Opt., 46, 4533, doi: 10.1364/AO.46.004533
- Dalcanton, J., Seager, S., Aigrain, S., et al. (2015), *From Cosmic Birth to Living Earths: The Future of UVOIR Space Astronomy*, AURA Report. <https://arxiv.org/abs/arXiv:1507.04779>
- Davies, R., Alves, J., Clénet, Y., et al. (2018), *The MICADO first light imager for the ELT: overview, operation, simulation*, Proc. SPIE, 10702, 107021S, doi: 10.1117/12.2311483
- Davies, R., Hörmann, V., Rabien, S., et al. (2021), *MICADO: The Multi-Adaptive Optics Camera for Deep Observations*, The ESO Messenger vol. 182, pp. 17-21, doi: 10.18727/0722-6691/5217
- Dean, B. H., Aronstein, D. L., Smith, J. S., Shiri, R., & Acton, D. S. (2006), *Phase retrieval algorithm for JWST Flight and Testbed Telescope*, Proc. SPIE, 6265, 626511, doi: 10.1117/12.673569
- Debes, J. H., Ren, B., & Schneider, G. (2019), *Pushing the Limits of the Coronagraphic Occulters on Hubble Space Telescope/Space Telescope Imaging Spectrograph*, JATIS, 5, 035003, doi: 10.1117/1.JATIS.5.3.035003
- Debes, J. H., Ygouf, M., Choquet, É., et al. (2015), *Wide-Field Infrared Survey Telescope–Astrophysics Focused Telescope Assets coronagraphic operations: lessons learned from the Hubble Space Telescope and the James Webb Space Telescope*, JATIS, 2, 1, doi: 10.1117/1.JATIS.2.1.011010
- Defrère, D., Hinz, P., Skemer, A., et al. (2015), *Exoplanet science with the LBTI: instrument status and plans*, Proc. SPIE, 9605, 96051G, doi: 10.1117/12.2188912
- Delorme, J. R., N'Diaye, M., Galicher, R., et al. (2016), *Laboratory Validation of the Dual-Zone Phase Mask Coronagraph in Broadband Light at the High-Contrast Imaging THD Testbed*, A&A, 592, A119, doi: 10.1051/0004-6361/201628587
- Desidera, S., Chauvin, G., Bonavita, M., et al. (2021), *The SPHERE Infrared Survey for Exoplanets (SHINE) - I. Sample Definition and Target Characterization*, A&A, 651, A70, doi: 10.1051/0004-6361/202038806
- Douglas, E. S., Carlton, A. K., Cahoy, K. L., et al. (2018), *WFIRST coronagraph technology requirements: status update and systems engineering approach*, Proc. SPIE, 10705, 1070526, doi: 10.1117/12.2314221
- Douglas, E. S., Males, J. R., Clark, J., et al. (2019), *Laser Guide Star for Large Segmented-Aperture Space Telescopes. I. Implications for Terrestrial Exoplanet Detection and Observatory Stability*, AJ, 157, 36, doi: 10.3847/1538-3881/aaf385

- East, M., Redding, D., Sullivan, C., Mooney, T., & Allen, L. (2018), *Picometer level stability of a mounted mirror assembly*, Proc. SPIE, 10698, 106980Y, doi: 10.1117/12.2314332
- East, M., Sullivan, C., & Wells, C. (2019), *Ultrastable mirror assembly design*, Proc. SPIE, 11116, 111160Z, Conference Presentation, doi: 10.1117/12.2529575
- Egron, S., Lajoie, C.-P., Leboulleux, L., et al. (2016), *James Webb Space Telescope optical simulation testbed III: first experimental results with linear-control alignment*, Proc. SPIE, 9904, 99044A, doi: 10.1117/12.2233650
- Egron, S., Lajoie, C.-P., Michau, V., et al. (2017a), *James Webb Space telescope optical simulation testbed: experimental results with linear control alignment*, Proc. SPIE, 10562, 1056225, doi: 10.1117/12.2296105
- Egron, S., Soummer, R., Lajoie, C.-P., et al. (2017b), *James Webb Space Telescope optical simulation testbed IV: linear control alignment of the primary segmented mirror*, Proc. SPIE, 10398, 1039811, doi: 10.1117/12.2272981
- Fanson, J., McCarthy, P. J., Bernstein, R., et al. (2018), *Overview and status of the Giant Magellan Telescope project*, Proc. SPIE, 10700, 1070012, doi: 10.1117/12.2313340
- Feinberg, L., Bolcar, M., Knight, S., & Redding, D. (2017), *Ultra-stable segmented telescope sensing and control architecture*, Proc. SPIE, 10398, 103980E, doi: 10.1117/12.2272810
- Fogarty, K., Mawet, D., Mazoyer, J., et al. (2020), *Towards high throughput and low-order aberration robustness for vortex coronagraphs with central obstructions*, Proc. SPIE, 11443, 114433Y, doi: 10.1117/12.2563180
- Foo, G., Palacios, D. M., & Swartzlander, G. A. (2005), *Optical Vortex Coronagraph*, Opt. Lett., 30, 3308, doi: 10.1364/OL.30.003308
- Galicher, R., Baudoz, P., & Rousset, G. (2008), *Wavefront Error Correction and Earth-like Planet Detection by a Self-Coherent Camera in Space*, A&A, 488, L9, doi: 10.1051/0004-6361:200810356
- Galicher, R., Huby, E., Baudoz, P., & Dupuis, O. (2020), *A Family of Phase Masks for Broadband Coronagraphy Example of the Wrapped Vortex Phase Mask Theory and Laboratory Demonstration*, A&A, 635, A11, doi: 10.1051/0004-6361/201936971
- Gaudi, B. S., Seager, S., Mennesson, B., et al. (2019), *The Habitable Exoplanet Observatory (HabEx) Mission Concept Study Final Report*. <https://arxiv.org/abs/arXiv:2001.06683>
- Give'on, A., Belikov, R., Shaklan, S., & Kasdin, J. (2007), *Closed Loop, DM Diversity-Based, Wavefront Correction Algorithm for High Contrast Imaging Systems*, Opt. Express, 15, 12338, doi: 10.1364/OE.15.012338
- Give'on, A., Kern, B. D., & Shaklan, S. (2011), *Pair-wise, deformable mirror, image plane-based diversity electric field estimation for high contrast coronagraphy*, Proc. SPIE, 8151, 815110, doi: 10.1117/12.895117
- Glindemann, A., Algomedo, J., Amestica, R., et al. (2003), *The VLTI: a status report*, Proc. SPIE, 4838, doi: 10.1117/12.458117
- Gonsalves, R. A. (1982), *Phase Retrieval And Diversity In Adaptive Optics*, Opt. Eng., 21, 215829, doi: 10.1117/12.7972989
- Grady, C. A., Proffitt, C. R., Malumuth, E., et al. (2003), *Coronagraphic Imaging with the Hubble Space Telescope and the Space Telescope Imaging Spectrograph1*, PASP, 115, 1036, doi: 10.1086/377083

- GRAVITY Collaboration, Abuter, R., Accardo, M., et al. (2017), *First light for GRAVITY: Phase referencing optical interferometry for the Very Large Telescope Interferometer*, A&A, 602, A94, doi: 10.1051/0004-6361/201730838
- Green, J. J. & Shaklan, S. B. (2003), *Optimizing coronagraph designs to minimize their contrast sensitivity to low-order optical aberrations*, Proc. SPIE, 5170, doi: 10.1117/12.506485
- Groff, T. D. & Kasdin, N. J. (2013), *Kalman filtering techniques for focal plane electric field estimation*, J. Opt. Soc. Am. A, 30, 128, doi: 10.1364/JOSAA.30.000128
- Groff, T. D., Riggs, A. J. E., Kern, B., & Kasdin, N. J. (2016), *Methods and Limitations of Focal Plane Sensing, Estimation, and Control in High-Contrast Imaging*, JATIS, 2, 011009, doi: 10.1117/1.JATIS.2.1.011009
- Guerri, G., Daban, J.-B., Robbe-Dubois, S., et al. (2011), *Apodized Lyot Coronagraph for SPHERE/VLT: II. Laboratory Tests and Performance*, Exp. Astron., 30, 59, doi: 10.1007/s10686-011-9220-y
- Guyon, O. (2003), *Phase-Induced Amplitude Apodization of Telescope Pupils for Extrasolar Terrestrial Planet Imaging*, A&A, 404, 379, doi: 10.1051/0004-6361:20030457
- Guyon, O. (2004), *Imaging Faint Sources within a Speckle Halo with Synchronous Interferometric Speckle Subtraction*, ApJ, 615, 562, doi: 10.1086/423980
- Guyon, O., Hinz, P. M., Cady, E., Belikov, R., & Martinache, F. (2014), *High Performance Lyot and PIAA Coronagraphy for Arbitrarily Shaped Telescope Apertures*, ApJ, 780, 18, doi: 10.1088/0004-637X/780/2/171
- Guyon, O., Lozi, J., Vievard, S., et al. (2020), *Validating advanced wavefront control techniques on the SCExAO testbed/instrument*, Proc. SPIE, 11448, 114481Z, doi: 10.1117/12.2562723
- Guyon, O., Martinache, F., Belikov, R., & Soummer, R. (2010a), *High Performance PIAA coronagraphy with Complex Amplitude Focal Plane Masks*, ApJS, 190, 220, doi: 10.1088/0067-0049/190/2/220
- Guyon, O., Martinache, F., Cady, E. J., et al. (2012), *How ELTs will acquire the first spectra of rocky habitable planets*, Proc. SPIE, 8447, 84471X, doi: 10.1117/12.927181
- Guyon, O., Martinache, F., Garrel, V., et al. (2010b), *The Subaru coronagraphic extreme AO (SCExAO) system: wavefront control and detection of exoplanets with coherent light modulation in the focal plane*, Proc. SPIE, 7736, 773624, doi: 10.1117/12.857878
- Guyon, O., Pluzhnik, E. A., Kuchner, M. J., Collins, B., & Ridgway, S. T. (2006), *Theoretical Limits on Extrasolar Terrestrial Planet Detection with Coronagraphs*, ApJ, 167, 81, doi: 10.1086/507630
- Hallibert, P., Boquet, F., Deslaef, N., & Sechi, G. (2019), *Technologies for large ultra-stable optical missions: current perspectives and developments at ESA*, Proc. SPIE, 11116, 1111610, doi: 10.1117/12.2529320
- Hartmann, J. (1900a), *Remarks on the Construction and Adjustment of Spectrographs. I.*, ApJ, 11, 400, doi: 10.1086/140709
- Hartmann, J. (1900b), *Remarks on the Construction and Adjustment of Spectrographs. II.*, ApJ, 12, 30, doi: 10.1086/140723
- Helmbrecht, M. A., He, M., & Kempf, C. J. (2013), *High-actuator-count MEMS deformable mirrors*, Proc. SPIE, 8725, 87250V, doi: 10.1117/12.2016198
- Helmbrecht, M. A., He, M., Kempf, C. J., & Marchis, F. (2016), *Long-term stability and temperature variability of Iris AO segmented MEMS deformable mirrors*, Proc. SPIE, 9909, 990981, doi: 10.1117/12.2233923

- Helmbrecht, M. A. & Juneau, T. (2007), *Piston-tip-tilt positioning of a segmented MEMS deformable mirror*, Proc. SPIE, 6467, 64670M, doi: 10.1117/12.702653
- Henry, G. W., Marcy, G., Butler, R. P., & Vogt, S. S. (1999), *HD 209458*, IAU Circ., 7307, 1
- Henry, G. W., Marcy, G. W., Butler, R. P., & Vogt, S. S. (2000), *A Transiting “51 Peg-like” Planet*, ApJL, 529, L41, doi: 10.1086/312458
- Herscovici-Schiller, O. (2018), *Analyse et correction de surface d’onde post-coronographique pour l’imagerie d’exoplanètes*, PhD thesis, Observatoire de Paris/Université Paris Sciences et Lettres
- Herscovici-Schiller, O., Mugnier, L. M., Baudoz, P., et al. (2018), *Towards the experimental validation of the non-linear dark hole on the THD bench*, Proc. SPIE, 10703, 1070329, doi: 10.1117/12.2314100
- Hinz, P., Bailey, V. P., Defrère, D., et al. (2014), *Commissioning the LBTI for use as a nulling interferometer and coherent imager*, Proc. SPIE, 9146, 91460T, doi: 10.1117/12.2057340
- Hinz, P., Codona, J., Guyon, O., et al. (2012), *TIGER: a high contrast infrared imager for the Giant Magellan Telescope*, Proc. SPIE, 8446, 84461P, doi: 10.1117/12.926751
- Hinz, P. M., Defrère, D., Skemer, A., et al. (2016), *Overview of LBTI: a multipurpose facility for high spatial resolution observations*, Proc. SPIE, 9907, 990704, doi: 10.1117/12.2233795
- Hodapp, K. W., Jensen, J. B., Irwin, E. M., et al. (2003), *The Gemini Near-Infrared Imager (NIRI)*, PASP, 115, 1388, doi: 10.1086/379669
- Houllé, M., Vigan, A., Carlotti, A., et al. (2021), *Direct Imaging and Spectroscopy of Exoplanets with the ELT/HARMONI High-Contrast Module*, A&A, 652, A67, doi: 10.1051/0004-6361/202140479
- Hunter, J. D. (2007), *Matplotlib: A 2D Graphics Environment*, Computing in Science and Engineering, 9, 90, doi: 10.1109/MCSE.2007.55
- Itoh, S., Matsuo, T., Shibai, H., & Sumi, T. (2019), *Point Spread Function of Hexagonally Segmented Telescopes by New Symmetrical Formulation*, MNRAS, 483, 119, doi: 10.1093/mnras/sty3052
- Jensen-Clem, R., Dillon, D., Gerard, B., et al. (2021), *The Santa Cruz Extreme AO Lab (SEAL): design and first light*, Proc. SPIE, 11823, 118231D, doi: 10.1117/12.2594676
- Jovanovic, N., Absil, O., Baudoz, P., et al. (2018), *Review of high-contrast imaging systems for current and future ground-based and space-based telescopes: Part II. Common path wavefront sensing/control and coherent differential imaging*, Proc. SPIE, 10703, 107031U, doi: 10.1117/12.2314260
- Jovanovic, N., Delorme, J.-R., Bond, C. Z., et al. (2019), *The Keck Planet Imager and Characterizer: demonstrating advanced exoplanet characterization techniques for future extremely large telescopes*, Proc. SPIE, 11117, 111170T, Conference Presentation, doi: 10.1117/12.2529330
- Jovanovic, N., Guyon, O., Lozi, J., et al. (2016), *The SCExAO high contrast imager: transitioning from commissioning to science*, Proc. SPIE, 9909, 99090W, doi: 10.1117/12.2234294
- Juanola-Parramon, R., Zimmerman, N., Groff, T., et al. (2019a), *Evaluating the LUVOIR Coronagraph Sensitivity to Telescope Aberrations*, IEEE Aerospace Conference Proceedings, doi: 10.1109/AERO.2019.8741658
- Juanola-Parramon, R., Zimmerman, N. T., Pueyo, L., et al. (2019b), *The LUVOIR Extreme Coronagraph for Living Planetary Systems (ECLIPS) II. Performance evaluation, aberration sensitivity analysis and exoplanet detection simulations*, Proc. SPIE, 11117, 1111702, doi: 10.1117/12.2530356
- Kalas, P., Graham, J. R., Fitzgerald, M. P., & Clampin, M. (2013), *STIS Coronagraphic Imaging of Fomalhaut: Main Belt Structure and the Orbit of Fomalhaut b*, ApJ, 775, 56, doi: 10.1088/0004-637X/775/1/56

- Kasdin, N. J., Bailey, V. P., Mennesson, B., et al. (2020), *The Nancy Grace Roman Space Telescope Coronagraph Instrument (CGI) technology demonstration*, Proc. SPIE, 11443, 114431U, doi: 10.1117/12.2562997
- Kasdin, N. J., Vanderbei, R. J., Spergel, D. N., & Littman, M. G. (2003), *Extrasolar Planet Finding via Optimal Apodized-Pupil and Shaped-Pupil Coronagraphs*, ApJ, 582, 1147, doi: 10.1086/344751
- Kasper, M., Cerpa Urra, N., Pathak, P., et al. (2021), *PCS — A Roadmap for Exoearth Imaging with the ELT*, The ESO Messenger vol. 182, doi: 10.18727/0722-6691/5221
- Kasper, M. E., Beuzit, J.-L., Verinaud, C., et al. (2008), *EPICS: the exoplanet imager for the E-ELT*, Proc. SPIE, 7015, 70151S, doi: 10.1117/12.789047
- Kenworthy, M. A., Mellon, S. N., Bailey, J. I., et al. (2021), *The β Pictoris b Hill Sphere Transit Campaign - I. Photometric Limits to Dust and Rings*, A&A, 648, A15, doi: 10.1051/0004-6361/202040060
- Knight, J. S., Acton, D. S., Lightsey, P., & Barto, A. (2012a), *Integrated telescope model for the James Webb Space Telescope*, Proc. SPIE, 8449, 84490V, doi: 10.1117/12.926814
- Knight, J. S., Acton, D. S., Lightsey, P., Contos, A., & Barto, A. (2012b), *Observatory alignment of the James Webb Space Telescope*, Proc. SPIE, 8442, 84422C, doi: 10.1117/12.926816
- Kopparapu, R. K. (2018), *The Habitable Zone: The Climatic Limits of Habitability* (Cham: Springer International Publishing), 2981–2993, doi: 10.1007/978-3-319-55333-7_58
- Kraus, S., Monnier, J. D., Ireland, M. J., et al. (2016), *Planet Formation Imager (PFI): science vision and key requirements*, Proc. SPIE, 9907, 99071K, doi: 10.1117/12.2231067
- Krist, J., Martin, S., Kuan, G., et al. (2019), *Numerical modeling of the Habex coronagraph*, Proc. SPIE, 11117, 1111705, doi: 10.1117/12.2530462
- Krist, J., Nemati, B., Zhou, H., & Sidick, E. (2015), *An overview of WFIRST/AFTA coronagraph optical modeling*, Proc. SPIE, 9605, 960505, doi: 10.1117/12.2188361
- Krist, J. E. (2004), *High-contrast imaging with the Hubble Space Telescope: performance and lessons learned*, Proc. SPIE, 5487, doi: 10.1117/12.548890
- Krist, J. E., Balasubramanian, K., Muller, R. E., et al. (2010), *The JWST/NIRCam coronagraph flight occulters*, Proc. SPIE, 7731, 77313J, doi: 10.1117/12.856488
- Krist, J. E., Beichman, C. A., Trauger, J. T., et al. (2007), *Hunting planets and observing disks with the JWST NIRCam coronagraph*, Proc. SPIE, 6693, 66930H, doi: 10.1117/12.734873
- Krist, J. E., Hartig, G. F., Clampin, M., et al. (2003), *Advanced camera for surveys coronagraph on the Hubble Space Telescope*, Proc. SPIE, 4860, doi: 10.1117/12.457835
- Kuhn, J. R., Potter, D., & Parise, B. (2001), *Imaging Polarimetric Observations of a New Circumstellar Disk System*, ApJ, 553, L189, doi: 10.1086/320686
- Lafrenière, D., Marois, C., Doyon, R., & Barman, T. (2009), *HST/NICMOS Detection of HR 8799 b in 1998*, ApJ, 694, L148, doi: 10.1088/0004-637X/694/2/L148
- Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, É. (2007), *A New Algorithm for Point-Spread Function Subtraction in High-Contrast Imaging: A Demonstration with Angular Differential Imaging*, ApJ, 660, 770, doi: 10.1086/513180
- Lagage, P. O. & Pantin, E. (1994), *Dust Depletion in the Inner Disk of β Pictoris as a Possible Indicator of Planets*, Nature, 369, 628, doi: 10.1038/369628a0
- Laginja, I. (2020), *PASTIS: v2.0.0 Semi-analytical PASTIS for LUVVOIR*, Zenodo, doi: 10.5281/zenodo.4170252

- Laginja, I. (2021), *spacetelescope/PASTIS*, Zenodo, doi: 10.5281/zenodo.3382985
- Laginja, I., Soummer, R., Mugnier, L. M., et al. (2021), *Analytical Tolerancing of Segmented Telescope Cophasing for Exo-Earth High-Contrast Imaging*, JATIS, 7, 015004, doi: 10.1117/1.JATIS.7.1.015004
- Laginja, I., Brady, G., Soummer, R., et al. (2018), *James Webb Space Telescope optical simulation testbed V: wide-field phase retrieval assessment*, Proc. SPIE, 10698, 106983N, doi: 10.1117/12.2313892
- Laginja, I., Lebouilleux, L., Pueyo, L., et al. (2019), *Wavefront error tolerancing for direct imaging of exo-Earths with a large segmented telescope in space*, Proc. SPIE, 11117, 1111717, doi: 10.1117/12.2530300
- Laginja, I., Soummer, R., Mugnier, L. M., et al. (2020), *Predicting contrast sensitivity to segmented aperture misalignment modes for the HiCAT testbed*, Proc. SPIE, 11443, 114433J, doi: 10.1117/12.2560113
- Laginja, I., Sauvage, J.-F., Mugnier, L. M., et al. (2022), *Wavefront tolerances of space-based segmented telescopes at very high contrast: Experimental validation*, A&A, 658, A84, doi: 10.1051/0004-6361/202142150
- Lagrange, A.-M., Bonnefoy, M., Chauvin, G., et al. (2010), *A Giant Planet Imaged in the Disk of the Young Star β Pictoris*, Science, 329, 57, doi: 10.1126/science.1187187
- Lagrange, A.-M., Kasper, M., Boccaletti, A., et al. (2009), *Constraining the Orbit of the Possible Companion to β Pictoris. New Deep Imaging Observations*, A&A, 506, 927, doi: 10.1051/0004-6361/200912098
- Lajoie, C.-P., Perrin, M. D., Myers, C., et al. (2018), *Wavefront sensing and control demo during the cryo-vacuum testing of JWST: exercising the science and operations center*, Proc. SPIE, 10698, 106983T, doi: 10.1117/12.2312145
- Lamb, M. P., Correia, C., Sivanandam, S., Swanson, R., & Zavyalova, P. (2021), *Simultaneous Estimation of Segmented Telescope Phasing Errors and Non-Common Path Aberrations from Adaptive-Optics-Corrected Images*, MNRAS, 505, 3347, doi: 10.1093/mnras/stab1247
- Langlois, M., Gratton, R., Lagrange, A.-M., et al. (2021), *The SPHERE Infrared Survey for Exoplanets (SHINE) - II. Observations, Data Reduction and Analysis, Detection Performances, and Initial Results*, A&A, 651, A71, doi: 10.1051/0004-6361/202039753
- Langlois, M., Vigan, A., Dohlen, K., et al. (2014), *Infrared differential imager and spectrograph for SPHERE: performance assessment for on-sky operation*, Proc. SPIE, 9147, 91479P, doi: 10.1117/12.2055535
- Langlois, M., Vigan, A., Moutou, C., et al. (2013), *Infrared Differential Imager and Spectrograph for SPHERE: Performance Status with Extreme Adaptive Optics before Shipment to ESO/VLT, AO4ELT3*, doi: 10.12839/AO4ELT3.13317
- Lebouilleux, L. (2018), *Optimal Wavefront Control for High-Contrast Imaging Application to Cophasing of Segmented Mirrors*, PhD thesis, Aix-Marseille Université
- Lebouilleux, L., Carlotti, A., & N'Diaye, M. (2021), *Redundant apodization for direct imaging of exoplanets 1: robustness to primary mirror segmentation-induced errors*, A&A, Submitted
- Lebouilleux, L., N'Diaye, M., Mazoyer, J., et al. (2017a), *Comparison of wavefront control algorithms and first results on the high-contrast imager for complex aperture telescopes (HiCAT) testbed*, Proc. SPIE, 10562, 105622Z, doi: 10.1117/12.2296154
- Lebouilleux, L., N'Diaye, M., Riggs, A. J. E., et al. (2016), *High-contrast imager for Complex Aperture Telescopes (HiCAT). 4. Status and wavefront control development*, Proc. SPIE, 9904, 99043C, doi: 10.1117/12.2233640

- Leboulleux, L., Pueyo, L., Sauvage, J.-F., et al. (2018a), *Sensitivity analysis for high-contrast imaging with segmented space telescopes*, Proc. SPIE, 10698, 106986H, doi: 10.1117/12.2313904
- Leboulleux, L., Sauvage, J.-F., Pueyo, L., et al. (2017b), *Sensitivity analysis for high-contrast missions with segmented telescopes*, Proc. SPIE, 10400, 104000M, doi: 10.1117/12.2274347
- Leboulleux, L., Sauvage, J.-F., Pueyo, L. A., et al. (2018b), *Pair-Based Analytical Model for Segmented Telescopes Imaging from Space for Sensitivity Analysis*, JATIS, 4, 15, doi: 10.1117/1.jatis.4.3.035002
- Leboulleux, L., Sauvage, J.-F., Soummer, R., et al. (2020), *Experimental Validation of Coronagraphic Focal-Plane Wavefront Sensing for Future Segmented Space Telescopes*, A&A, 639, A70, doi: 10.1051/0004-6361/202037658
- Lenzen, R., Hartung, M., Brandner, W., et al. (2003), *NAOS-CONICA first on-sky results in a variety of observing modes*, Proc. SPIE, 4841, doi: 10.1117/12.460044
- Lightsey, P. A., Atkinson, C. B., Clampin, M. C., & Feinberg, L. D. (2012), *James Webb Space Telescope: Large Deployable Cryogenic Telescope in Space*, Opt. Eng., 51, 011003, doi: 10.1117/1.OE.51.1.011003
- Lightsey, P. A. & Chrisp, M. (2003), *Image quality for large segmented space telescopes*, Proc. SPIE, 4850, doi: 10.1117/12.461220
- Lightsey, P. A., Knight, J. S., Bolcar, M. R., et al. (2018), *Optical budgeting for LUVVOIR*, Proc. SPIE, 10698, 1069813, doi: 10.1117/12.2312256
- Lightsey, P. A., Knight, J. S., & Golnik, G. (2014), *Status of the optical performance for the James Webb Space Telescope*, Proc. SPIE, 9143, 914304, doi: 10.1117/12.2055502
- Lissauer, J. J. (1987), *Timescales for Planetary Accretion and the Structure of the Protoplanetary Disk*, ICARUS, 69, 249, doi: 10.1016/0019-1035(87)90104-7
- Llop-Sayson, J., Ruane, G., Mawet, D., et al. (2020), *High-contrast Demonstration of an Apodized Vortex Coronagraph*, AJ, 159, 79, doi: 10.3847/1538-3881/ab6329
- Llop-Sayson, J., Ruane, G., Mawet, D. P., et al. (2019), *Demonstration of an Electric Field Conjugation Algorithm for Improved Starlight Rejection through a Single Mode Optical Fiber*, JATIS, 5, 019004, doi: 10.1117/1.JATIS.5.1.019004
- Lou, J. Z., Redding, D. C., Nissen, J. A., & Shelton, C. (2018), *LUVVOIR primary mirror segment alignment control with joint laser metrology and segment edge sensing*, Proc. SPIE, 10698, 1069840, doi: 10.1117/12.2311495
- Lozi, J., Guyon, O., Jovanovic, N., et al. (2018), *SCEXAO, an instrument with a dual purpose: perform cutting-edge science and develop new technologies*, Proc. SPIE, 10703, 1070359, doi: 10.1117/12.2314282
- Lumbres, J., Males, J., Douglas, E., et al. (2018), *Modeling coronagraphic extreme wavefront control systems for high contrast imaging in ground and space telescope missions*, Proc. SPIE, 10703, 107034Z, doi: 10.1117/12.2313780
- Lyon, R. G. & Clampin, M. (2012), *Space Telescope Sensitivity and Controls for Exoplanet Imaging*, Opt. Eng., 51, 16, doi: 10.1117/1.OE.51.1.011002
- Lyot, B. (1932), *Étude de La Couronne Solaire En Dehors Des Éclipses. Avec 16 Figures Dans Le Texte.*, Zeitschrift für Astrophysik, 5, 73
- Lyot, B. (1939), *The Study of the Solar Corona and Prominences without Eclipses (George Darwin Lecture, Delivered by M. Bernard Lyot, Assoc.R.A.S., on 1939 May 12)*, MNRAS, 99, 580, doi: 10.1093/mnras/99.8.580

- Macintosh, B., Chilcote, J. K., Bailey, V. P., et al. (2018), *The Gemini Planet Imager: looking back over five years and forward to the future*, Proc. SPIE, 10703, 107030K, doi: 10.1117/12.2314253
- Macintosh, B., Graham, J. R., Barman, T., et al. (2015), *Discovery and Spectroscopy of the Young Jovian Planet 51 Eri b with the Gemini Planet Imager*, Science, 350, 64, doi: 10.1126/science.aac5891
- Macintosh, B. A., Anthony, A., Atwood, J., et al. (2014), *The Gemini planet imager: first light and commissioning*, Proc. SPIE, 9148, 91480J, doi: 10.1117/12.2056709
- Malbet, F., Yu, J. W., & Shao, M. (1995), *High-Dynamic-Range Imaging Using a Deformable Mirror for Space Coronagraphy*, PASP, 107, 386, doi: 10.1086/133563
- Males, J. R., Belikov, R., & Bendek, E. (2015), *Orbital Differential Imaging: a new high-contrast post-processing technique for direct imaging of exoplanets*, Proc. SPIE, 9605, 960518, doi: 10.1117/12.2188766
- Males, J. R., Close, L. M., Guyon, O., et al. (2020), *MagAO-X first light*, Proc. SPIE, 11448, 114484L, doi: 10.1117/12.2561682
- Males, J. R., Close, L. M., Miller, K., et al. (2018), *MagAO-X: project status and first laboratory results*, Proc. SPIE, 10703, 1070309, doi: 10.1117/12.2312992
- Mamajek, E. E. & Bell, C. P. M. (2014), *On the Age of the β Pictoris Moving Group*, MNRAS, 445, 2169, doi: 10.1093/mnras/stu1894
- Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. (2006), *Angular Differential Imaging: A Powerful High-Contrast Imaging Technique*, ApJ, 641, 556, doi: 10.1086/500401
- Marois, C., Macintosh, B., Barman, T., et al. (2008), *Direct Imaging of Multiple Planets Orbiting the Star HR 8799*, Science, 322, 1348, doi: 10.1126/science.1166585
- Marois, C., Macintosh, B., & Véran, J.-P. (2010a), *Exoplanet imaging with LOCI processing: photometry and astrometry with the new SOSIE pipeline*, Proc. SPIE, 773, 77361J, doi: 10.1117/12.857225
- Marois, C., Zuckerman, B., Konopacky, Q. M., Macintosh, B., & Barman, T. (2010b), *Images of a Fourth Planet Orbiting HR 8799*, Nature, 468, 1080, doi: 10.1038/nature09684
- Martinez, P., Boccaletti, A., Kasper, M., et al. (2008), *Comparison of Coronagraphs for High-Contrast Imaging in the Context of Extremely Large Telescopes*, A&A, 492, 289, doi: 10.1051/0004-6361:200810650
- Mast, T. S., Nelson, J. E., & Welch, W. J. (1982), *Effects Of Primary Mirror Segmentation On Telescope Image Quality*, Proc. SPIE, 0332, doi: 10.1117/12.933512
- Matthews, B., Kennedy, G., Sibthorpe, B., et al. (2013), *Resolved Imaging of the HR 8799 Debris Disk with Herschel*, ApJ, 780, 97, doi: 10.1088/0004-637X/780/1/97
- Mawet, D., Pueyo, L., Carlotti, A., et al. (2013), *Ring-Apodized Vortex Coronagraphs for Obscured Telescopes. I. Transmissive Ring Apodizers*, ApJS, 209, 7, doi: 10.1088/0067-0049/209/1/7
- Mawet, D., Pueyo, L., Lawson, P., et al. (2012), *Review of small-angle coronagraphic techniques in the wake of ground-based second-generation adaptive optics systems*, Proc. SPIE, 8442, 844204, doi: 10.1117/12.927245
- Mawet, D., Riaud, P., Absil, O., & Surdej, J. (2005), *Annular Groove Phase Mask Coronagraph*, ApJ, 633, 1191, doi: 10.1086/462409
- Mawet, D., Wizinowich, P., Dekany, R., et al. (2016), *Keck Planet Imager and Characterizer: concept and phased implementation*, Proc. SPIE, 9909, 99090D, doi: 10.1117/12.2233658

- Mayor, M. & Queloz, D. (1995), *A Jupiter-Mass Companion to a Solar-Type Star*, *Nature*, 378, 355, doi: 10.1038/378355a0
- Mazoyer, J. (2014), *Haut Contraste Pour l'imagerie Directe d'exoplanètes et de Disques : De La Self-Coherent Camera à l'analyse de Données NICI*, PhD thesis, Université Paris Diderot – Paris 7
- Mazoyer, J., Baudoz, P., Belikov, R., et al. (2019), *High-Contrast Testbeds for Future Space-Based Direct Imaging Exoplanet Missions*, *Bulletin of the AAS*, 51, Retrieved from <https://baas.aas.org/pub/2020n7i101>
- Mazoyer, J., Baudoz, P., Galicher, R., Mas, M., & Rousset, G. (2013), *Estimation and Correction of Wavefront Aberrations Using the Self-Coherent Camera: Laboratory Results*, *A&A*, 557, 9, doi: 10.1051/0004-6361/201321706
- Mazoyer, J., Pueyo, L., N'Diaye, M., et al. (2018a), *Active Correction of Aperture Discontinuities-Optimized Stroke Minimization. I. A New Adaptive Interaction Matrix Algorithm*, *AJ*, 155, 13, doi: 10.3847/1538-3881/aa91cf
- Mazoyer, J., Pueyo, L., N'Diaye, M., et al. (2018b), *Active Correction of Aperture Discontinuities-Optimized Stroke Minimization. II. Optimization for Future Missions*, *AJ*, 155, 19, doi: 10.3847/1538-3881/aa91d7
- McKinney, W. (2010), *Data Structures for Statistical Computing in Python*, 9th Python in Science Conference, doi: 10.25080/Majora-92bf1922-00a
- Mennesson, B., Debes, J., Douglas, E., et al. (2018), *The WFIRST coronagraph instrument: a major step in the exploration of sun-like planetary systems via direct imaging*, *Proc. SPIE*, 10698, 106982I, doi: 10.1117/12.2313861
- Mocœur, I., Mugnier, L. M., & Cassaing, F. (2009), *Analytical Solution to the Phase-Diversity Problem for Real-Time Wavefront Sensing*, *Opt. Lett.*, 34, 3487, doi: 10.1364/OL.34.003487
- Moody, D. C. & Trauger, J. T. (2007), *Hybrid Lyot coronagraph masks and wavefront control for improved spectral bandwidth and throughput*, *Proc. SPIE*, 6693, 66931I, doi: 10.1117/12.735070
- Moore, D. B. & Redding, D. C. (2018), *Picometer differential wavefront metrology by nonlinear Zernike wavefront sensing for LUVVOIR*, *Proc. SPIE*, 10698, 106984I, doi: 10.1117/12.2312600
- Moriarty, C., Brooks, K., Soummer, R., et al. (2018), *High-contrast imager for complex aperture telescopes (HiCAT): 6. software control infrastructure and calibration*, *Proc. SPIE*, 10698, 1069853, doi: 10.1117/12.2314058
- Mouillet, D., Larwood, J. D., Papaloizou, J. C. B., & Lagrange, A. M. (1997), *A Planet on an Inclined Orbit as an Explanation of the Warp in the β Pictoris Disc*, *MNRAS*, 292, 896, doi: 10.1093/mnras/292.4.896
- Mugnier, L. M., Cornia, A., Sauvage, J.-F., et al. (2009), *Optimal Method for Exoplanet Detection by Angular Differential Imaging*, *JOSA A*, 26, 1326, doi: 10.1364/JOSAA.26.001326
- Navascués, D. B., Stauffer, J. R., Song, I., & Caillault, J.-P. (1999), *The Age of β Pictoris*, *ApJ*, 520, L123, doi: 10.1086/312162
- N'Diaye, M., Choquet, É., Egron, S., et al. (2014), *High-contrast Imager for Complex Aperture Telescopes (HiCAT): II. Design overview and first light results*, *Proc. SPIE*, 9143, 914327, doi: 10.1117/12.2056694
- N'Diaye, M., Choquet, É., Pueyo, L., et al. (2013), *High-contrast imager for complex aperture telescopes (HiCAT): 1. testbed design*, *Proc. SPIE*, 8864, 88641K, doi: 10.1117/12.2023718
- N'Diaye, M., Mazoyer, J., Choquet, É., et al. (2015a), *High-contrast imager for complex aperture telescopes (HiCAT): 3. first lab results with wavefront control*, *Proc. SPIE*, 9605, 96050I, doi: 10.1117/12.2188497

- N'Diaye, M., Pueyo, L., & Soummer, R. (2015b), *Apodized Pupil Lyot Coronagraphs for Arbitrary Apertures. IV. Reduced Inner Working Angle and Increased Robustness to Low-Order Aberrations*, *ApJ*, 799, 225, doi: 10.1088/0004-637X/799/2/225
- N'Diaye, M., Soummer, R., Pueyo, L., et al. (2016), *Apodized Pupil Lyot Coronagraphs for Arbitrary Apertures. V. Hybrid Shaped Pupil Designs for Imaging Earth-like Planets with Future Space Observatories*, *AJ*, 818, 163, doi: 10.3847/0004-637x/818/2/163
- Nemati, B., Krist, J. E., & Mennesson, B. (2017a), *Sensitivity of the WFIRST coronagraph performance to key instrument parameters*, *Proc. SPIE*, 10400, 1040007, doi: 10.1117/12.2274396
- Nemati, B., Stahl, H. P., Stahl, M. T., Ruane, G. J. J., & Sheldon, L. J. (2020), *Method for Deriving Optical Telescope Performance Specifications for Earth-Detecting Coronagraphs*, *JATIS*, 6, 039002, doi: 10.1117/1.JATIS.6.3.039002
- Nemati, B., Stahl, M. T., Stahl, H. P., & Shaklan, S. B. (2017b), *The effects of space telescope primary mirror segment errors on coronagraph instrument performance*, *Proc. SPIE*, 10398, 103980G, doi: 10.1117/12.2273072
- Nielsen, E. L., de Rosa, R. J., Macintosh, B., et al. (2019), *The Gemini Planet Imager Exoplanet Survey: Giant Planet and Brown Dwarf Demographics from 10 to 100 Au*, *AJ*, 158, 13, doi: 10.3847/1538-3881/ab16e9
- Noll, R. J. (1976), *Zernike polynomials and atmospheric turbulence*, *J. Opt. Soc. Am.*, 66, 207, doi: 10.1364/JOSA.66.000207
- Noss, J., Fowler, J., Moriarty, C., et al. (2021a), *spacetelescope/catkit*, Zenodo, doi: 10.5281/zenodo.4312171
- Noss, J., Fowler, J., Moriarty, C., et al. (2021b), *spacetelescope/catkit: v0.36.1*, Zenodo, Zenodo, doi:10.5281/zenodo.5076238, doi: 10.5281/zenodo.5076238
- Oliphant, T. E. (2006), *A guide to NumPy*, Vol. 1 (Trelgol Publishing USA), <https://web.mit.edu/dvp/Public/numpybook.pdf>
- Otten, G. P. P. L., Snik, F., Kenworthy, M. A., et al. (2017), *On-Sky Performance Analysis of the Vector Apodizing Phase Plate Coronagraph on MagAO/Clio2*, *ApJ*, 834, 175, doi: 10.3847/1538-4357/834/2/175
- Patru, F., Baudoz, P., Galicher, R., et al. (2018), *In-lab testing of six-level phase mask coronagraphs onto the high-contrast imaging THD2 bench*, *Proc. SPIE*, 10703, 107032L, doi: 10.1117/12.2314152
- Patterson, K. D., Seo, B.-J., Balasubramanian, K., et al. (2019), *Design description and commissioning performance of a stable coronagraph technology development testbed for direct imaging of Earth-like exoplanets*, *Proc. SPIE*, 11117, 111171U, doi: 10.1117/12.2530401
- Paul, B., Mugnier, L. M., Sauvage, J.-F., & Dohlen, K. (2013), *High-Order Myopic Coronagraphic Phase Diversity (COFFEE) for Wave-Front Control in High-Contrast Imaging Systems*, *Opt. Express*, 21, 31751, doi: 10.1364/OE.21.031751
- Paul, B., Sauvage, J.-F., Mugnier, L. M., et al. (2014a), *Simultaneous phase and amplitude retrieval with COFFEE: from theory to laboratory results*, *Proc. SPIE*, 9147, doi: 10.1117/12.2055447
- Paul, B., Sauvage, J.-F., Mugnier, L. M., et al. (2014b), *Compensation of High-Order Quasi-Static Aberrations on SPHERE with the Coronagraphic Phase Diversity (COFFEE)*, *A&A*, 572, A32, doi: 10.1051/0004-6361/201424133
- Perrin, M., Long, J., Douglas, E., et al. (2016a), *POPPY: Physical Optics Propagation in PYTHON*, ASCL. <http://ascl.net/ascl:1504.007>

- Perrin, M. D., Acton, D. S., Lajoie, C.-P., et al. (2016b), *Preparing for JWST wavefront sensing and control operations*, Proc. SPIE, 9904, 99040F, doi: 10 . 1117/12 . 2233104
- Perrin, M. D., Pueyo, L., van Gorkom, K., et al. (2018), *Updated optical modeling of JWST coronagraph performance contrast, stability, and strategies*, Proc. SPIE, 10698, 1069809, doi: 10 . 1117/12 . 2313552
- Perrin, M. D., Sivaramakrishnan, A., Lajoie, C.-P., et al. (2014a), *Updated point spread function simulations for JWST with WebbPSF*, Proc. SPIE, 9143, 91433X, doi: 10 . 1117/12 . 2056689
- Perrin, M. D., Soummer, R., Choquet, É., et al. (2014b), *James Webb Space Telescope Optical Simulation Testbed I: overview and first results*, Proc. SPIE, 9143, 914309, doi: 10 . 1117/12 . 2056936
- Perrin, M. D., Soummer, R., Elliott, E. M., Lallo, M. D., & Sivaramakrishnan, A. (2012), *Simulating point spread functions for the James Webb Space Telescope with WebbPSF*, Proc. SPIE, 8442, 84423D, doi: 10 . 1117/12 . 925230
- Pogorelyuk, L. & Kasdin, N. J. (2019), *Dark Hole Maintenance and A Posteriori Intensity Estimation in the Presence of Speckle Drift in a High-contrast Space Coronagraph*, ApJ, 873, 95, doi: 10 . 3847/1538-4357/ab0461
- Pogorelyuk, L., Pueyo, L., Males, J. R., Cahoy, K., & Kasdin, N. J. (2021), *Information-Theoretical Limits of Recursive Estimation and Closed-Loop Control in High-Contrast Imaging*, ApJS, 256, 39, doi: 10 . 3847/1538-4365/ac126d
- Polidan, R. S., Belvin, W. K., Greenhouse, M. A., et al. (2018), *Servicing and assembly: enabling the most ambitious future space observatories*, Proc. SPIE, 10698, 1069825, doi: 10 . 1117/12 . 2313013
- Por, E. H. (2020a), *Novel approaches for direct exoplanet imaging: Theory, simulations and experiments*, PhD thesis, Leiden Observatory, Faculty of Science, Leiden University
- Por, E. H. (2020b), *Phase-Apodized-Pupil Lyot Coronagraphs for Arbitrary Telescope Pupils*, ApJ, 888, 127, doi: 10 . 3847/1538-4357/ab3857
- Por, E. H., Haffert, S. Y., Radhakrishnan, V. M., et al. (2018), *High Contrast Imaging for Python (HCIPy): an open-source adaptive optics and coronagraph simulator*, Proc. SPIE, 10703, 1070342, doi: 10 . 1117/12 . 2314407
- Por, E. H., Soummer, R., Noss, J., & St.Laurent, K. (2020), *Exploiting symmetries and progressive refinement for apodized pupil Lyot coronagraph design*, Proc. SPIE, 11443, 114433P, doi: 10 . 1117/12 . 2562571
- Potier, A., Baudoz, P., Galicher, R., Singh, G., & Boccaletti, A. (2020a), *Comparing Focal Plane Wavefront Control Techniques: Numerical Simulations and Laboratory Experiments*, A&A, 635, A192, doi: 10 . 1051/0004-6361/201937015
- Potier, A., Galicher, R., Baudoz, P., et al. (2020b), *Increasing the Raw Contrast of VLT/SPHERE with the Dark Hole Technique - I. Simulations and Validation on the Internal Source*, A&A, 638, A117, doi: 10 . 1051/0004-6361/202038010
- Potier, A., Ruane, G., Chen, P., et al. (2021), *LUVOIR-ECLIPS closed-loop adaptive optics performance and contrast predictions*, Proc. SPIE, 11823, 118231L, doi: 10 . 1117/12 . 2595116
- Pourcelot, R., N'Diaye, M., Brady, G., et al. (2020), *Estimating low-order aberrations through a Lyot coronagraph with a Zernike wavefront sensor for exoplanet imaging*, Proc. SPIE, 11443, doi: 10 . 1117/12 . 2562159
- Pourcelot, R., N'Diaye, M., Por, E. H., et al. (2021), *Experimental validation of active control of low-order aberrations with a Zernike sensor through a Lyot coronagraph*, Proc. SPIE, 11823, doi: 10 . 1117/12 . 2594609

- Pourcelot, R., N'Diaye, M., Por, R. S. E. H., et al. (2022), *Low-order wavefront control using a Zernike sensor through Lyot coronagraphs for exoplanet imaging: Blind stabilization of an image dark hole*, A&A, Submitted
- Pueyo, L., Crepp, J. R., Vasisht, G., et al. (2012), *Application of a Damped Locally Optimized Combination of Images Method to the Spectral Characterization of Faint Companions Using an Integral Field Spectrograph*, ApJS, 199, 6, doi: 10.1088/0067-0049/199/1/6
- Pueyo, L., Juanola-Parramon, R., Tumlinson, J., et al. (2021), *Coronagraphic detection of earth-like planets with large, actively-controlled space telescopes*, JATIS, In preparation
- Pueyo, L., Kay, J., Kasdin, N. J., et al. (2009), *Optimal Dark Hole Generation via Two Deformable Mirrors with Stroke Minimization*, Appl. Opt., 48, 17, doi: 10.1364/AO.48.006296
- Pueyo, L., Soummer, R., Knight, J. S., & Coyle, I. (2019), *Stability error budget for exo-earth imaging with a large segmented telescope in space.*, AAS Meeting, 233, Conference poster, <https://aas233-aas.ipostersessions.com/default.aspx?s=F3-CD-DE-E7-D7-CD-24-BC-8B-57-9C-0B-DA-73-A1-00&guestview=true>
- Pueyo, L., Stark, C., Juanola-Parramon, R., et al. (2019), *The LUVOIR Extreme Coronagraph for Living Planetary Systems (ECLIPS) I: searching and characterizing exoplanetary gems*, Proc. SPIE, 11117, 1111703, doi: 10.1117/12.2530722
- Pueyo, L., Zimmerman, N., Bolcar, M., et al. (2017), *The LUVOIR architecture "A" coronagraph instrument*, Proc. SPIE, 10398, 103980F, doi: 10.1117/12.2274654
- Pueyo, L. A., Pogorelyuk, L., Laginja, I., et al. (2021), *Fundamental limits of high contrast imaging with continuous WFS&C control in space.*, Proc. SPIE, 11823, 1182345, Conference Presentation, doi: 10.1117/12.2594955
- Racine, R., Walker, G. A. H., Nadeau, D., Doyon, R., & Marois, C. (1999), *Speckle Noise and the Detection of Faint Companions*, PASP, 111, 587, doi: 10.1086/316367
- Ragazzoni, R. (1996), *Pupil Plane Wavefront Sensing with an Oscillating Prism*, J. Mod. Opt., 43, 289, doi: 10.1080/09500349608232742
- Rajan, A., Barman, T., Soummer, R., et al. (2015), *Characterizing the Atmospheres of the HR 8799 Planets with HST/WFC3*, ApJ, 809, L33, doi: 10.1088/2041-8205/809/2/L33
- Rajan, A., Rameau, J., de Rosa, R. J., et al. (2017), *Characterizing 51 Eri b from 1 to 5 μ m: A Partly Cloudy Exoplanet*, AJ, 154, 10, doi: 10.3847/1538-3881/aa74db
- Rameau, J., Chauvin, G., Lagrange, A.-M., et al. (2015), *Detection Limits with Spectral Differential Imaging Data*, A&A, 581, A80, doi: 10.1051/0004-6361/201525879
- Ramsay, S., Amico, P., Bezawada, N., et al. (2020), *The ESO Extremely Large Telescope instrumentation programme*, Proc. SPIE, 11203, 1120303, doi: 10.1117/12.2541400
- Reback, J., McKinney, W., jbrockmendel, et al. (2020), *pandas-dev/pandas: Pandas 1.1.4*, Zenodo, doi: 10.5281/zenodo.4161697
- Redmond, S. F., Pogorelyuk, L., Pueyo, L., et al. (2022), *Implementation of a dark zone maintenance algorithm for speckle drift correction in a high contrast space coronagraph*, JATIS, Submitted
- Redmond, S. F., Pueyo, L., Pogorelyuk, L., et al. (2021), *Dark zone maintenance results for segmented aperture wavefront error drift in a high contrast space coronagraph*, Proc. SPIE, 11823, 118231K, doi: 10.1117/12.2594647

- Redmond, S. M., Kasdin, N. J., Pogorelyuk, L., et al. (2020), *Implementation of a dark hole maintenance algorithm for speckle drift in a high contrast space coronagraph*, Proc. SPIE, 11443, doi: 10.1117/12.2561488
- Rencher, A. C. & Schaalje, G. B. (2008), *Linear models in statistics*, 2nd edn. (Wiley-Interscience), doi: 10.1002/9780470192610
- Richichi, A. & Percheron, I. (2005), *First Results from the ESO VLTI Calibrators Program*, A&A, 434, 1201, doi: 10.1051/0004-6361:20042257
- Riggs, A. E., Ruane, G., & Kern, B. D. (2019), *Directly constraining low-order aberration sensitivities in the WFIRST coronagraph design*, Proc. SPIE, 11117, 111170F, doi: 10.1117/12.2529588
- Riggs, A. J. E., Bailey, V., Moody, D. C., et al. (2021), *Flight mask designs of the Roman Space Telescope coronagraph instrument*, Proc. SPIE, 11823, 118231Y, doi: 10.1117/12.2598599
- Riggs, A. J. E., Kasdin, N. J., & Groff, T. D. (2016), *Recursive starlight and bias estimation for high-contrast imaging with an extended Kalman filter*, JATIS, 2, 1, doi: 10.1117/1.JATIS.2.1.011017
- Riggs, A. J. E., Ruane, G., Fogarty, K., Pueyo, L., & Balasubramanian, K. (2018), *Numerically optimized coronagraph designs for the Habitable Exoplanet Imaging Mission (HabEx) concept*, Proc. SPIE, 10698, 106980V, doi: 10.1117/12.2313843
- Roberge, A., Bolcar, M. R., & France, K. C. (2019), *Telling the story of life in the cosmos: the LUVOIR telescope concepts*, Proc. SPIE, 11115, 111150O, doi: 10.1117/12.2530475
- Roddier, F. (1999), *Adaptive Optics in Astronomy* (Cambridge University Press), doi: 10.1017/CBO9780511525179
- Roddier, F. & Roddier, C. (1997), *Stellar Coronagraph with Phase Mask*, PASP, 109, 815, doi: 10.1086/133949
- Roman, N. G. (1959), *Planets of Other Suns.*, AJ, 64, 344, doi: 10.1086/108038
- Roques, F., Scholl, H., Sicardy, B., & Smith, B. A. (1994), *Is There a Planet around β Pictoris? Perturbations of a Planet on a Circumstellar Dust Disk: 1. The Numerical Model*, ICARUS, 108, 37, doi: 10.1006/icar.1994.1040
- Rouan, D., Riaud, P., Boccaletti, A., Clénet, Y., & Labeyrie, A. (2000), *The Four-Quadrant Phase-Mask Coronagraph. I. Principle*, PASP, 112, 1479, doi: 10.1086/317707
- Rousset, G., Lacombe, F., Puget, P., et al. (2003), *NAOS, the first AO system of the VLT: on-sky performance*, Proc. SPIE, 4839, doi: 10.1117/12.459332
- Ruane, G., Crill, B., Patterson, K., et al. (2019), *Decadal Survey Testbed Commissioning Roadmap : Demonstrating Technology for Imaging New Worlds*, Tech. rep., NASA ExEP, https://exoplanets.nasa.gov/internal_resources/1170/
- Ruane, G., Mawet, D., Jewell, J., & Shaklan, S. (2017), *Performance and sensitivity of vortex coronagraphs on segmented space telescopes*, Proc. SPIE, 10400, 104000J, doi: 10.1117/12.2274508
- Ruane, G., Mawet, D., Mennesson, B., Jewell, J., & Shaklan, S. (2018a), *Vortex Coronagraphs for the Habitable Exoplanet Imaging Mission Concept: Theoretical Performance and Telescope Requirements*, JATIS, 4, 1, doi: 10.1117/1.jatis.4.1.015004
- Ruane, G., Riggs, A., Mazoyer, J., et al. (2018b), *Review of high-contrast imaging systems for current and future ground- and space-based telescopes I: coronagraph design methods and optical performance metrics*, Proc. SPIE, 10698, 106982S, doi: 10.1117/12.2312948

- Ruane, G. J., Wallace, J. K., Steeves, J. B., et al. (2020), *Wavefront Sensing and Control in Space-Based Coronagraph Instruments Using Zernike's Phase-Contrast Method*, JATIS, 6, 045005, doi: 10.1117/1.JATIS.6.4.045005
- Sahoo, A., Laginja, I., Pueyo, L. A., et al. (2021), *LUVOIR A Segment level thermo-mechanical stability requirements*, Proc. SPIE, 11823, 1182349, Conference Presentation, doi: 10.1117/12.2594636
- Sahoo, A., Pueyo, L., Laginja, I., et al. (2022), *LUVOIR A segment-level thermal stability analysis*, JATIS, In preparation
- Samland, M., Mollière, P., Bonnefoy, M., et al. (2017), *Spectral and Atmospheric Characterization of 51 Eridani b Using VLT/SPHERE*, A&A, 603, A57, doi: 10.1051/0004-6361/201629767
- Sauvage, J.-F., Mugnier, L., Fusco, T., & Rousset, G. (2006), *Post processing of differential images for direct extrasolar planet detection from the ground*, Proc. SPIE, 6272, 62722B, doi: 10.1117/12.671470
- Sauvage, J.-F., Mugnier, L., Paul, B., & Villicroze, R. (2012), *Coronagraphic Phase Diversity: A Simple Focal Plane Sensor for High-Contrast Imaging*, Opt. Lett., 37, 4808, doi: 10.1364/OL.37.004808
- Sauvage, J.-F., Mugnier, L. M., Rousset, G., & Fusco, T. (2010), *Analytical expression of long-exposure adaptive-optics-corrected coronagraphic image. First application to exoplanet detection*, J. Opt. Soc. Am. A, 27, A157, doi: 10.1364/JOSAA.27.00A157
- Savransky, D. & Garrett, D. (2015), *WFIRST-AFTA Coronagraph Science Yield Modeling with EXOSIMS*, JATIS, 2, 011006, doi: 10.1117/1.jatis.2.1.011006
- Schmid, H. M., Bazzon, A., Roelfsema, R., et al. (2018), *SPHERE/ZIMPOL High Resolution Polarimetric Imager - I. System Overview, PSF Parameters, Coronagraphy, and Polarimetry*, A&A, 619, A9, doi: 10.1051/0004-6361/201833620
- Schneider, G., Debes, J. H., Grady, C. A., et al. (2018), *The HR 4796A Debris System: Discovery of Extensive Exo-Ring Dust Material*, AJ, 155, 77, doi: 10.3847/1538-3881/aaa3f3
- Schneider, G., Grady, C. A., Hines, D. C., et al. (2014), *Probing for Exoplanets Hiding in Dusty Debris Disks: Disk Imaging, Characterization, and Explorations with HST/STIS Multi-Roll Coronagraphy*, AJ, 148, 59, doi: 10.1088/0004-6256/148/4/59
- Seo, B.-J., Cady, E., Gordon, B., et al. (2017), *Hybrid Lyot coronagraph for WFIRST: high-contrast broadband testbed demonstration*, Proc. SPIE, 10400, 104000F, doi: 10.1117/12.2274687
- Seo, B.-J., Patterson, K., Balasubramanian, K., et al. (2019), *Testbed demonstration of high-contrast coronagraph imaging in search for Earth-like exoplanets*, Proc. SPIE, 11117, 111171V, doi: 10.1117/12.2530033
- Shack, R. V., Rancourt, J. D., & Morrow, H. (1971), *Effects of Dilution on a Six-Element Synthetic Aperture*, Appl. Opt., 10, 257, doi: 10.1364/AO.10.000257
- Shaklan, S. B. & Green, J. J. (2005), *Low-Order Aberration Sensitivity of Eighth-Order Coronagraph Masks*, ApJ, 628, 474, doi: 10.1086/430204
- Shaklan, S. B., Green, J. J., & Palacios, D. M. (2006), *The Terrestrial Planet Finder Coronagraph optical surface requirements*, Proc. SPIE, 6265, 62651I, doi: 10.1117/12.670124
- Shaklan, S. B., Marchen, L., Green, J. J., & Lay, O. P. (2005), *The Terrestrial Planet Finder Coronagraph dynamics error budget*, Proc. SPIE, 5905, 59050D, doi: 10.1117/12.617890
- Shectman, S. A. & Johns, M. (2003), *The Magellan Telescopes*, Proc. SPIE, 4837, doi: 10.1117/12.457909

- Shi, F., Ohara, C. M., Chanan, G., Troy, M., & Redding, D. C. (2004), *Experimental verification of dispersed fringe sensing as a segment-phasing technique using the Keck Telescope*, Proc. SPIE, 5489, doi: 10.1117/12.552407
- Sidick, E. & Riggs, A. E. (2019), *Optimization of Lyot coronagraph focal plane masks for improved low-order wavefront sensing*, Proc. SPIE, 11117, 111170G, doi: 10.1117/12.2529876
- Sidick, E., Shaklan, S., Balasubramanian, K., & Cady, E. (2015), *Studies of the effects of control bandwidth and dark-hole size on the HCIT contrast performance*, Proc. SPIE, 9605, 96050H, doi: 10.1117/12.2188210
- Silvester, S., Tanbakuchi, A., Müller, P., et al. (2020), *imageio/imageio v2.8.0*, Zenodo, doi: 10.5281/zenodo.3674133
- Simard, L., Ellerbroek, B., Bhatia, R., Radovan, M., & Chisholm, E. (2016), *Thirty Meter Telescope science instruments: a status report*, Proc. SPIE, 9908, 99081V, doi: 10.1117/12.2234058
- Sivaramakrishnan, A., Soummer, R., Oppenheimer, B. R., et al. (2010), *Gemini Planet Imager coronagraph testbed results*, Proc. SPIE, 7735, 773586, doi: 10.1117/12.858623
- Sivaramakrishnan, A. & Yaitskova, N. (2005), *Lyot Coronagraphy on Giant Segmented-Mirror Telescopes*, ApJ, 626, L65, doi: 10.1086/431460
- Smith, B. A. & Terrile, R. J. (1984), *A Circumstellar Disk around Beta Pictoris*, Science, 226, 1421, doi: 10.1126/science.226.4681.1421
- Smith, W. H. (1987), *Spectral Differential Imaging Detection of Planets about Nearby Stars.*, PASP, 99, 1344, doi: 10.1086/132124
- Snik, F., Otten, G., Kenworthy, M., et al. (2012), *The Vector-APP: a Broadband Apodizing Phase Plate that yields Complementary PSFs*, Proc. SPIE, 8450, 84500M, doi: 10.1117/12.926222
- Soummer, R. (2004), *Apodized Pupil Lyot Coronagraphs for Arbitrary Telescope Apertures*, ApJ, 618, L161, doi: 10.1086/427923
- Soummer, R., Aime, C., & Falloon, P. E. (2003a), *Stellar Coronagraphy with Prolate Apodized Circular Apertures*, A&A, 397, 1161, doi: 10.1051/0004-6361:20021573
- Soummer, R., Brady, G. R., Brooks, K., et al. (2018), *High-contrast imager for complex aperture telescopes (HiCAT): 5. first results with segmented-aperture coronagraph and wavefront control*, Proc. SPIE, 10698, 106981O, doi: 10.1117/12.2314110
- Soummer, R., Dohlen, K., & Aime, C. (2003b), *Achromatic Dual-Zone Phase Mask Stellar Coronagraph*, A&A, 403, 369, doi: 10.1051/0004-6361:20030278
- Soummer, R., Hagan, J. B., Pueyo, L., et al. (2011), *Orbital Motion of HR 8799 b, c, d Using Hubble Space Telescope Data from 1998: Constraints on Inclination, Eccentricity, and Stability*, ApJ, 741, 55, doi: 10.1088/0004-637x/741/1/55
- Soummer, R., Pueyo, L., Ferrari, A., et al. (2009), *Apodized Pupil Lyot Coronagraphs for Arbitrary Apertures. II. Theoretical Properties and Application to Extremely Large Telescopes*, ApJ, 695, 695, doi: 10.1088/0004-637x/695/1/695
- Soummer, R., Pueyo, L., & Larkin, J. (2012), *Detection and Characterization of Exoplanets and Disks Using Projections on Karhunen-Loève Eigenimages*, ApJ, 755, L28, doi: 10.1088/2041-8205/755/2/L28
- Stahl, H. P. (2017), *Advanced mirror technology development (AMTD): year five status*, Proc. SPIE, 10401, 104010O, doi: 10.1117/12.2275190

- Stahl, H. P., Kuan, G., Arnold, W. R., et al. (2020), *Habitable-Zone Exoplanet Observatory Baseline 4-m Telescope: Systems-Engineering Design Process and Predicted Structural Thermal Optical Performance*, JATIS, 6, 034004, doi: 10.1117/1.JATIS.6.3.034004
- Stahl, M. T., Shaklan, S. B., & Stahl, H. P. (2015), *Preliminary analysis of effect of random segment errors on coronagraph performance*, Proc. SPIE, 9605, 96050P, doi: 10.1117/12.2190160
- Stark, C. C., Belikov, R., Bolcar, M. R., et al. (2019), *ExoEarth yield landscape for future direct imaging space telescopes*, JATIS, 5, 1, doi: 10.1117/1.JATIS.5.2.024009
- Stark, C. C., Roberge, A., Mandell, A., et al. (2015), *Lower Limits on Aperture Size for an ExoEarth Detecting Coronagraphic Mission*, ApJ, 808, 149, doi: 10.1088/0004-637x/808/2/149
- Steeves, J., Wallace, J. K., Kettenbeil, C., & Jewell, J. (2020), *Picometer Wavefront Sensing Using the Phase-Contrast Technique*, Optica, 7, 1267, doi: 10.1364/OPTICA.398768
- Su, K. Y. L., Rieke, G. H., Stapelfeldt, K. R., et al. (2009), *The Debris Disk Around HR 8799*, ApJ, 705, 314, doi: 10.1088/0004-637x/705/1/314
- Thalmann, C., Schmid, H. M., Boccaletti, A., et al. (2008), *SPHERE ZIMPOL: overview and performance simulation*, Proc. SPIE, 7014, 70143F, doi: 10.1117/12.789158
- Thatte, N., Tecza, M., Schnetler, H., et al. (2021), *HARMONI: The ELT's First-Light Near-Infrared and Visible Integral Field Spectrograph*, The ESO Messenger vol. 182, pp. 7-12, doi: 10.18727/0722-6691/5215
- The Astropy Collaboration (2018), *astropy: a core python package for astronomy*, Zenodo, doi: 10.5281/zenodo.1461536
- The Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. (2018), *The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package*, AJ, 156, 123, doi: 10.3847/1538-3881/aabc4f
- The Astropy Collaboration, Robitaille, Thomas P., Tollerud, Erik J., et al. (2013), *Astropy: A community Python package for astronomy*, A&A, 558, A33, doi: 10.1051/0004-6361/201322068
- The LUVOIR Team (2019), *The LUVOIR Mission Concept Study Final Report*. <https://arxiv.org/abs/arXiv:1912.06219>
- Trauger, J., Moody, D., Gordon, B., Krist, J., & Mawet, D. (2012), *Complex apodization Lyot coronagraphy for the direct imaging of exoplanet systems: design, fabrication, and laboratory demonstration*, Proc. SPIE, 8442, 84424Q, doi: 10.1117/12.926663
- Trauger, J. T., Burrows, C., Gordon, B., et al. (2004), *Coronagraph contrast demonstrations with the high-contrast imaging testbed*, Proc. SPIE, 5487, doi: 10.1117/12.552520
- Trauger, J. T. & Traub, W. A. (2007), *A Laboratory Demonstration of the Capability to Image an Earth-like Extrasolar Planet*, Nature, 446, 771, doi: 10.1038/nature05729
- Troy, M. & Chanan, G. A. (2003), *Diffraction effects from giant segmented mirror telescopes*, Proc. SPIE, 4840, doi: 10.1117/12.457985
- van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., et al. (2014), *scikit-image: Image processing in Python*, PeerJ, doi: 10.7717/peerj.453
- Vigan, A., Fontanive, C., Meyer, M., et al. (2021), *The SPHERE infrared survey for exoplanets (SHINE) - III. The demographics of young giant exoplanets below 300 au with SPHERE*, A&A, 651, A72, doi: 10.1051/0004-6361/202038107

- Vilas, F. & Smith, B. A. (1987), *Coronagraph for astronomical imaging and spectrophotometry*, Appl. Opt., 26, 664, doi: 10.1364/AO.26.000664
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. (2020), *SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python*, Nature Methods, Vol. 17, pp. 261-272, 17, 261, doi: 10.1038/s41592-019-0686-2
- Wagner, K., Apai, D., & Kratter, K. M. (2019), *On the Mass Function, Multiplicity, and Origins of Wide-Orbit Giant Planets*, ApJ, 877, 46, doi: 10.3847/1538-4357/ab1904
- Wang, J. J., Graham, J. R., Dawson, R., et al. (2018), *Dynamical Constraints on the HR 8799 Planets with GPI*, AJ, 156, 192, doi: 10.3847/1538-3881/aae150
- Wang, J. J., Graham, J. R., Pueyo, L., et al. (2016), *The Orbit and Transit Prospects for β Pictoris b Constrained with One Milliarcsecond Astrometry*, AJ, 152, 97, doi: 10.3847/0004-6256/152/4/97
- Wells, C. & East, M. (2019), *Mirror segment sensitivity analysis and performance budgets for LUVUOIR ULTRA*, Proc. SPIE, 11116, 1111603, Conference Presentation, doi: 10.1117/12.2529405
- Will, S. D. (2021), *Advanced coronagraphic modeling and wavefront control*, PhD thesis, The Institute of Optics, University of Rochester
- Will, S. D., Groff, T. D., & Fienup, J. R. (2021a), *Jacobian-Free Coronagraphic Wavefront Control Using Non-linear Optimization*, JATIS, 7, doi: 10.1117/1.JATIS.7.1.019002
- Will, S. D., Perrin, M. D., Por, E. H., et al. (2022), *High-order coronagraphic wavefront control with algorithmic differentiation: first experimental demonstration*, JATIS, In preparation
- Will, S. D., Perrin, M. D., Por, E. H., et al. (2021b), *Wavefront control with algorithmic differentiation on the HiCAT testbed*, Proc. SPIE, 11823, doi: 10.1117/12.2594283
- Wizinowich, P., Acton, D. S., Shelton, C., et al. (2000), *First Light Adaptive Optics Images from the Keck II Telescope: A New Era of High Angular Resolution Imagery*, PASP, 112, 315, doi: 10.1086/316543
- Yaitskova, N. & Dohlen, K. (2002), *Tip-tilt error for extremely large segmented telescopes: detailed theoretical point-spread-function analysis and numerical simulation results*, J. Opt. Soc. Am. A, 19, 1274, doi: 10.1364/JOSAA.19.001274
- Yaitskova, N., Dohlen, K., & Dierickx, P. (2003), *Analytical study of diffraction effects in extremely large segmented telescopes*, J. Opt. Soc. Am. A, 20, 1563, doi: 10.1364/JOSAA.20.001563
- Yaitskova, N. & Troy, M. (2011), *Rolled edges and phasing of segmented telescopes*, Appl. Opt., 50, 542, doi: 10.1364/AO.50.000542
- Ygouf, M., Mugnier, L. M., Mouillet, D., Fusco, T., & Beuzit, J.-L. (2013), *Simultaneous Exoplanet Detection and Instrument Aberration Retrieval in Multispectral Coronagraphic Imaging*, A&A, 551, A138, doi: 10.1051/0004-6361/201220318
- Zernike, F. & Stratton, F. (1934), *Diffraction Theory of the Knife-Edge Test and Its Improved Form, The Phase-Contrast Method*, MNRAS, 94, 377, doi: 10.1093/mnras/94.5.377
- Zimmerman, N. T., N'Diaye, M., St-Laurent, K. E., et al. (2016), *Lyot coronagraph design study for large, segmented space telescope apertures*, Proc. SPIE, 9904, 99041Y, doi: 10.1117/12.2233205

ACRONYMS

ACE	Ames Coronagraph Experiment
ADI	angular differential imaging
ANDROMEDA	ANgular DiffeRential OptiMal Exoplanet Detection Algorithm
AO	adaptive optics
API	application programming interface
APLC	apodized pupil Lyot coronagraph
AVC	apodizer vortex coronagraph
CA	core accretion
CDI	coherent differential imaging
CGI	Coronagraph Instrumnet
CL	closed-loop
CLC	classical Lyot coronagraph
COFFEE	COronagraphic Focal-plane wave-Front Estimation for Exoplanet detection
DH	dark hole
DM	deformable mirror
DMVC	deformable-mirror-supported vortex coronagraph
dOTF	differential optical transfer function
DST	Decadal Survey Testbed
DZ	dark zone
E2E	end-to-end
ECLIPS	Extreme Coronagraph for Living Planetary Systems
EFC	electric field conjugation
ELASTIC	Estimation of Large Amplitude Subaperture Tip-tilt by Image Correlation
ELT	Extremely Large Telescope
ESA	European Space Agency
ESO	European Southern Observatory
ExEP	Exoplanet Exploration Program

FP	focal plane
FPM	focal-plane mask
FQPM	four-quadrant phase mask [coronagraph]
GI	gravitational instability
GMT	Giant Magellan Telescope
GPCT	General Purpose Coronagraph Testbed
GPI	Gemini Planet Imager
GPIES	Gemini Planet Imager Exoplanet Survey
GPR	Geometric Phase Retrieval
GRD	Groupe Recherche et Développement
HabEx	Habitable Exoplanet Observatory
HARMONI	High Angular Resolution Monolithic Optical and Near-infrared Integral field spectrograph
HCG	HabEx Coronagraph
HCI	high-contrast imaging
HCIT	High-Contrast Imaging Testbed
HCST	High-Contrast Spectroscopy Testbed for Segmented Telescopes
HDA	Hybrid Diversity Algorithm
HiCAT	High-contrast imager for Complex Aperture Telescopes [testbed]
HLC	hybrid Lyot coronagraph
HST	Hubble Space Telescope
IFS	integral field spectrograph
IR	infrared
ITM	Integrated Telescope Model
IWA	inner working angle
JOST	JWST Optical Simulation Testbed
JPL	Jet Propulsion Laboratory
JWST	James Webb Space Telescope
KPIC	Keck Planet Imager and Characterizer
L2	Lagrange point 2
LAM	Laboratoire d'Astrophysique de Marseille
LAPD	linearized analytic phase diversity
LBT	Large Binocular Telescope

LBTI	Large Binocular Telescope Interferometer
LOCI	locally optimized combination of images
LOWFS	low-order wavefront sensor
LS	Lyot stop
LSB	least significant bit
LUVOIR	Large UltraViolet Optical InfraRed Surveyor
MagAO-X	Magellan telescope extreme adaptive optics
MAORY	Multi-conjugate Adaptive Optics RelaY
MC	Monte Carlo
METIS	Mid-infrared ELT Imager and Spectrograph
MICADO	Multi-AO Imaging Camera for Deep Observations
MIRI	Mid-Infrared Instrument
NASA	National Aeronautics and Space Administration
NIRCam	Near Infrared Camera
OAP	off-axis parabola
OBWFS	out-of-band wavefront sensor
ODI	orbital differential imaging
OL	open-loop
OMC	Occulting Mask Coronagraph [testbed]
ONERA	Office National d'Etudes et Recherches Aérospatiales
OOP	object-oriented programming
OPERA	Optimized Phase Retrieval Algorithm
OTE	Optical Telescope Element
OTF	optical transfer function
OWA	outer working angle
PAPLC	phase-apodized pupil Lyot coronagraph
PASTIS	Pair-based Analytical model for Segmented Telescope Imaging from Space
PCS	Planetary Camera and Spectrograph
PDI	polarization differential imaging
PFI	Planet Formation Imager
PIAA	phase-induced amplitude apodization [coronagraph]
PIAACMC	phase-induced amplitude apodization complex mask coronagraph

PP	pupil plane
PSD	power spectral density
PSF	point spread function
PTT	piston-tip-tilt
PW	pair-wise [estimator]
R&R	Roddier and Roddier [coronagraph]
RMOL	Russell B. Makidon Optics Laboratory
RST	Roman Space Telescope
RV	radial velocity
SCC	self-coherent camera
SCDA	segmented coronagraph design analysis
SCExAO	Subaru Coronagraphic Extreme AO [system]
SDI	spectral differential imaging
SEAL	Santa Cruz Extreme AO Lab
SHINE	SpHere INfrared survey for Exoplanets
SM	stroke minimization
SNR	signal-to-noise ratio
SOC	Science and Operations Center
SPC	shaped pupil coronagraph
SPEED	Segmented Pupil Experiment for Exoplanet Detection
SPHERE	Spectro-Polarimetric High-contrast Exoplanet REsearch
SSI	Starshade Instrument
STScI	Space Telescope Science Institute
TBT	Testbed Telescope
THD2	Très Haute Dynamique 2 [testbed]
TMA	three-mirror anastigmat
TMT	Thirty Meter Telescope
UV	ultraviolet
UVOIR	ultraviolet/optical/infrared
vAPP	vector apodizing phase plate [coronagraph]
VC	vortex coronagraph
VLT	Very Large Telescope

VLTI	Very Large Telescope Interferometer
WF	wavefront
WFE	wavefront error
WFS	wavefront sensor
WFS&C	wavefront sensing and control
xAO	extreme adaptive optics
ZWFS	Zernike wavefront sensor

ENGLISH SUMMARY

For centuries, humans have wondered and imagined what lies beyond our home planet, the Earth. Early philosophers asked themselves whether planets orbiting stars other than our Sun could exist. Concrete scientific theories addressing this question emerged in the 19th and early 20th centuries, but in 1995, the first detection of a planet orbiting a star that is similar to our Sun led us to seeing these extrasolar planets, *exoplanets*, in a completely new way. Not only are there many more planetary systems out there than just our own, but the diversity in planets has been found to be very different compared to what we know about the bodies in our solar system. With so many foreign worlds out there, the question arises: is there other life out there too?

Detecting exoplanets requires sophisticated instruments that capture the signals of these hard-to-detect objects: since it is easier to observe the much brighter host star rather than a planet directly, the first detection methods that flourished and unveiled thousands of new bodies are what we call *indirect* detection methods. These do not collect the light of the planet directly, instead they measure effects that let us infer the existence of one or more planets in orbit around a star. The two most fruitful methods in this sense are planetary *transits*¹ and the *radial velocity* method². When an exoplanet transits its host star through our line of sight toward that star, we can detect the dimming this causes in the star light, from which we can determine the exoplanet's existence and its radius. The radial velocity detection method observes a periodic shift in the wavelength of the light that reaches us from the star, which is caused by its back-and-forth motion from the gravitational tug the planet exerts on it; this also tells us an approximate mass of the exoplanet.

Thousands of exoplanets have been detected with such indirect methods, but to be able to characterize these distant worlds with more than just their radius and mass, we need to gain direct access to the light they reflect or emit themselves, rather than using the host star as a proxy. This is enabled by a *direct* detection method, which blocks the overwhelming light from the star and takes images of the planets themselves. This allows us to do a spectral analysis of their atmospheres, telling us what molecules they contain and whether they would be able to support alien life. This method is called *direct imaging*, and there are a couple of hurdles to overcome to be able to use it to detect exoplanets. First, when observed from a distance, these planets are very close to their host star, with projected separations in the sky of one arcsecond to several milli-arcseconds. Resolving the planet as an individual point source thus requires a good *angular resolution*, which is enabled by telescopes with very large mirrors. Second, exoplanets are very faint, and their light gets lost in the glare of their host star, much like trying to see a lightning bug right next to a bright streetlight. This *brightness ratio*, or *contrast*, depends on the wavelength we observe at: in the infrared, where young, giant planets still glow from their own formation process, this brightness ratio is around $10^{-5} - 10^{-6}$. In visible wavelengths, where planets reflect the light from their host star, this ratio becomes even harder to overcome as it is as wide as 10^{-10} . Assuming that a planet needs to closely resemble the Earth to have a chance to support life, we need to be looking for rocky and relatively small planets, orbiting a main-sequence star like the Sun at a distance of one astronomical unit, which is the distance between the Sun and the Earth. An important goal of direct imaging is thus to find such “*exoEarths*”. The instruments capable of such observations will need to be able to reach a contrast of 10^{-10} in reflected light in the visible wavelength range, at an angular separation of 0.1 arcsec.

To overcome these two challenges that are imposed by the fundamental, physical behavior of light, astronomers use complex *high-contrast imaging* instruments that provide, as their name suggests, high-contrast observation capabilities at small angular resolution. There are three main components in such an

¹<https://exoplanets.nasa.gov/alien-worlds/ways-to-find-a-planet/#/2>

²<https://exoplanets.nasa.gov/alien-worlds/ways-to-find-a-planet/#/1>

instrument that each contribute to the overall imaging performance. The first of them, the *coronagraph*, is a set of optical masks that aim to reduce the overall glare of the star while preserving the light coming from the planet. The first concept for such a coronagraph was successfully built and used by French astronomer Bernard Lyot in 1930, when he used a combination of a *focal-plane mask* and what we today call a *Lyot stop* to block out the disk of the Sun in order to observe the faint solar corona. It was only decades later that scientists adapted such instruments to be pointed at other stars, allowing us to observe faint circumstellar disks - and planets.

The second important component of a high-contrast instrument arises from the distortions light is exposed to before it reaches the detector. While unobstructed light from a distant, point-like object travels as a plane wave, it can encounter sources of aberrations along the way that introduce a *wavefront error*. This can happen due to the Earth's atmosphere, in which temperature and pressure gradients, as well as wind, cause significant turbulence to deteriorate the quality of the final image significantly. Moreover, there are sources of wavefront errors internal to the telescope that also contribute to this effect, like misalignments of various components in the optical system, or polishing errors on optical surfaces. The component that mitigates these effects to a certain degree is a *wavefront sensing and control system*. It consists of a wavefront sensor that detects the aberrations introducing the errors, and control components like *deformable mirrors* that apply corrections by adapting their own surface shapes according to the measurement of the sensor. Together with the coronagraph, such control loops are used to modulate the light in a designated part of the image plane, creating a dark area of very high contrast, the *dark hole*.

The third big component of high-contrast imaging is post-processing. While it is preferable to separate the planet and star light with an appropriate optical system before they get detected, there are data reduction methods that allow us to improve the detected contrast in the final image, pushing the limit for detection even further.

With these enabling technologies, there has been an abundance of exciting science results, for example from circumstellar disk images like the ones shown in Fig. 1.2 in Chap. 1. The first exoplanet detection with direct imaging was done in 2004, and several others followed. Observations of warm, giant exoplanets like the ones shown in Fig. 1.3 are being used to constrain their planetary orbits and learn about their atmospheres. However, with the instruments we have today, we are not able to see small, rocky exoplanets that might potentially harbor life. Even with the 30 m-class ground-based telescopes coming online in the mid 2030s, also called extremely large telescopes, or ELTs, we will be limited to contrast levels that will at best reach 10^{-9} , precluding us from seeing Earth-like exoplanets around Sun-type stars. To open the window to those objects, we have to put our telescopes into space to avoid the adverse influence of the Earth's atmosphere; and we have to develop more powerful high-contrast imaging instruments than what we use today.

Sending a telescope capable of exoEarth imaging to space poses a significant set of challenges. To achieve an angular resolution and required sensitivity that are sufficient to image a faint point source that is only 0.1 arcsec away from a bright star, the observatory needs to contain a primary mirror diameter of several meters, which yields a very large light collecting area. The engineering solution to build large apertures of more than ~ 4 m in space is to compose them of individual, hexagonal segments, assembled to form the total primary mirror. This makes them lighter than monolithic equivalents, easier to manufacture, and they can also be folded up to fit into the fairing of the rocket that brings them into space. The first segmented telescope sent to space will be the James Webb Space Telescope, with 18 segments combined to give a mirror with a total diameter of 6.5 meters, and it will provide us with valuable learning opportunities for future, even more ambitious segmented telescopes. While a large telescope aperture gives us the necessary angular resolution and sensitivity, the overall architecture of the high-contrast imaging system defines what contrast levels we can achieve during observations.

There is a whole range of coronagraph designs that are suitable for direct imaging from space, but each of them comes with trade-offs in their respective capabilities: different *inner and outer working angles* (how close or far from the star they can observe), *planet throughput* (how much of the planet light they lose while eliminating the stellar light) or compatibility with segment gaps in the telescope pupil. However, they are all conceived to be able to reach 10^{-10} contrast under perfect conditions - without aberrations and without temporal evolution of the optical system. While there is no fast-evolving atmosphere in space that

introduces aberrations, there are still several significant sources of wavefront errors on a space telescope, which the coronagraph is very sensitive to. Imperfections in the manufacturing of the optical components and surfaces introduce distortions that an on-board wavefront sensing and control system can correct for to a certain extent. More aberrations are introduced by thermally and mechanically induced deformations of the whole observatory structure, which degrade the contrast. This is especially true for deformations originating from the backplane structure of the primary mirror, which is holding together the array of individual segments. Some of these wavefront errors dynamically evolve during the observations and the question arises whether we can keep such a large, segmented telescope stable enough so that the contrast does not degrade to the point where we become blind to exoEarths. The goal of this thesis is to develop methods to characterize the necessary limits for these wavefront error variations, in particular for the distortions introduced by primary mirror segments, and to prepare future space telescopes for the efficient use of high-contrast instruments in space.

Chapter 2: The theoretical basis of wavefront error tolerancing for large segmented telescopes

Misalignments between the segments of a large telescope contaminate the images of coronagraphic observations with blobs of light, so-called *speckles*, in exactly the areas where we want to observe exoplanets, the dark hole. Before we can decide how to mitigate this problem, it is important to characterize the exact effect of segment-level misalignments on the coronagraphic contrast. The PASTIS model for propagations of segmented aberrations through a coronagraph converts these wavefront errors truthfully into an average contrast number and provides an analytical basis that allows us to invert the problem. Instead of asking: “What is the contrast under x amount of aberrations?”, we ask: “What level x of wavefront errors can we tolerate if we need a specific contrast for our observation?” This chapter shows how to derive these wavefront error tolerances with analytical equations and introduces their statistical interpretation over a range of many different misalignment realizations of the segmented mirror. A major result that emerged from the work in this chapter is that not all segments on a segmented telescope carrying a coronagraph need to be constrained equally strictly.

Chapter 3: Numerical simulation tools and simulated tolerancing results

Numerical simulators are an indispensable tool for the demonstration of an idea on a carefully implemented model. The PASTIS tolerancing package is a software unit written in Python that performs a tolerancing analysis as presented in Chap. 2. It contains a range of optical simulators that reproduce the telescope geometry and coronagraph of several observatory designs, which can be connected to the calculations of wavefront error tolerances. It also allows us to connect simulators not contained in this package to perform the analysis. Applying these theoretical and numerical methods to the concrete case of a coronagraph on the James Webb Space Telescope gives us an idea of how the contrast performance of this instrument relates to potential misalignments on the primary mirror of this telescope. Aiming closer to home, this chapter also presents simulated tolerancing experiments on the HiCAT testbed.

Chapter 4: Experimental validations of segmented wavefront error tolerances on HiCAT

The HiCAT testbed at STScI is an optical laboratory experiment for the demonstration of technologies for high-contrast imaging from space. It contains a Lyot coronagraph, two continuous deformable mirrors for wavefront sensing and control, and a segmented deformable mirror that reproduces the optical effects of a large, segmented primary mirror. Its 37 segments can be controlled individually in piston, tip and tilt to introduce misalignments. This chapter shows the results from experiments with this setup in which the analytical equations developed in Chap. 2 are confirmed on a real high-contrast instrument. The conclusions from this work allow us to determine how well we need to control the segmented deformable mirror if we want to reach contrast levels that lie beyond the current performance of HiCAT.

Chapter 5: Keeping large, segmented telescopes in space stable enough for exoEarth imaging

The analytical tolerancing framework developed in Chap. 2 is in fact not limited to misalignment effects from segmented mirrors. The method is equally applicable to wavefront errors that span the entire telescope pupil at once, or to treat specific aberrations that arise from thermal and mechanical deformations, found through modeling. This opens a whole slew of applications for this very fundamental treatment of aberrations. This includes going a step further, looking more closely at how certain wavefront aberrations

evolve over time, and what would be appropriate measures to counteract the effect of evolving aberrations on the coronagraphic contrast. This chapter first presents some quantitative results from the tolerancing of thermally induced wavefront errors, and then showcases the role of such methods in the overall design of wavefront sensing and control strategies of large, segmented telescopes in space.

Appendix A: Demonstrating wavefront sensing and control on a segmented testbed

Before any observation with a segmented telescope can start, the segment misalignments need to be reduced by means of wavefront sensing and control. For observations that include a coronagraph, the segments must be aligned to a level equal to or better than required by the tolerancing analysis developed in Chap. 2. The segmented, but non-coronagraphic JOST testbed at STScI uses a wavefront sensing and control algorithm called LAPD to find the best alignment solution. This chapter presents the results of LAPD experiments on JOST on a wide field of view with the particular goal to align the segmented mirror of the testbed.

Appendix B: Coronagraphy on a segmented-aperture laboratory testbed

Many coronagraphs and wavefront sensing and control technologies have been developed to reach a high contrast on space telescopes. They usually work best in the perfect case of circular apertures that are not obstructed by secondary mirrors and spiders, and under perfect conditions when there are no wavefront errors present in the optical system. However, real observatories contain a whole range of features that make reaching a high-contrast DH hard, like the ones listed above. The goal of the HiCAT testbed is to investigate how to bring various components of high-contrast instruments like sensors, controllers and starlight suppression techniques together into an efficiently functioning system operating on obscured apertures that include segmentation. This chapter lists the main motivation and design concept of this testbed, and explains the various testbed setups that have been used to conduct wavefront sensing and control experiments with a segmented-aperture coronagraph.

The ultimate reason for the efforts put into instrumentation for high-contrast imaging is to improve the scientific results we can obtain with our telescopes. To learn about the existence, but also the overall occurrence of Earth-like exoplanets in our galactic neighborhood, what we seek to increase is the *exoplanet candidate yield*, meaning we need to find a large enough number of potentially habitable planets to start making more global conclusions, such as whether the Earth is a rather typical planet in our galaxy or not. While there is currently no space telescope that uses actively controlled coronagraphs, future missions will give us that capability. The monolithic 2.4 m Roman Space Telescope will be the first observatory in space to fly deformable mirrors as an intermediate milestone, and future missions like the segmented 15 m LUVOIR telescope concept will deploy the means for an exoEarth-imaging observatory.

I sincerely thank Eimear O'Reilly, Mia Mayer and Christopher Moriarty for proof-reading my thesis, and in particular this summary.

RÉSUMÉ EN FRANÇAIS

Depuis des siècles, les humains s'interrogent et imaginent ce qui se trouve au-delà de notre planète, la Terre. Les premiers philosophes se sont demandé s'il existait des planètes orbitant autour d'autres étoiles que notre Soleil. Si des théories scientifiques concrètes répondant à cette question ont vu le jour au 19e et au début du 20e siècle, la première détection d'une planète en orbite autour d'une étoile similaire à notre Soleil en 1995 nous a amené à voir ces planètes extrasolaires, les *exoplanètes*, d'une manière totalement nouvelle. Non seulement il existe beaucoup plus de systèmes planétaires que le nôtre, mais la diversité des planètes s'est avérée très différente de ce que nous connaissons des corps de notre système solaire. Avec un tel nombre de mondes extrasolaires, une question se pose : y a-t-il aussi de la vie ailleurs ?

La détection d'exoplanètes nécessite des instruments sophistiqués pour capter les signaux de ces objets difficiles à identifier : comme il est plus facile d'observer directement l'étoile hôte, beaucoup plus brillante, plutôt qu'une planète, les premières méthodes de détection qui ont émergé et dévoilé des milliers de nouveaux objets sont ce que nous appelons des méthodes de détection *indirectes*. Celles-ci n'enregistrent pas directement la lumière de la planète, mais mesurent des effets qui nous permettent de déduire l'existence d'une ou plusieurs planètes en orbite autour d'une étoile. Les deux méthodes les plus fructueuses dans ce sens sont les *transits* planétaires³ et les *vitesses radiales*⁴. Un transit se produit lorsqu'une exoplanète traverse notre ligne de visée vers son étoile hôte. Nous pouvons alors détecter l'atténuation provoquée dans la luminosité de l'étoile, à partir de laquelle nous pouvons déterminer l'existence et le rayon de l'exoplanète. La méthode de détection par vitesse radiale permet d'observer un décalage périodique de la longueur d'onde de la lumière qui nous parvient de l'étoile. Ce décalage est causé par le mouvement de va-et-vient de l'étoile induit par la force gravitationnelle que la planète exerce sur elle ; cela procure également une approximation de la masse de l'exoplanète.

Des milliers d'exoplanètes ont été détectées à l'aide de ces méthodes indirectes. Cependant, pour pouvoir caractériser ces mondes lointains autrement que par leur rayon et leur masse, nous devons avoir un accès direct à la lumière qu'ils reflètent ou émettent eux-mêmes, plutôt que d'utiliser l'étoile hôte comme intérimaire. Ceci est possible grâce à une méthode de détection *directe*, qui consiste à bloquer la lumière dominante de l'étoile et à imager les planètes elles-mêmes. Cela permet d'effectuer une analyse spectrale de leur atmosphère, nous indiquant leur composition moléculaire et si elles sont susceptibles d'accueillir une vie extraterrestre. Cette méthode, appelée *imagerie directe*, nécessite de surmonter quelques obstacles afin de l'utiliser pour détecter des exoplanètes. Premièrement, lorsqu'elles sont observées à distance, ces planètes sont très proches de leur étoile hôte, avec des séparations projetées dans le ciel d'une seconde d'arc à plusieurs millisecondes d'arc. La résolution de la planète en tant que source ponctuelle individuelle nécessite donc une haute *résolution angulaire*, permise par les télescopes dotés de très grands miroirs. Deuxièmement, les exoplanètes sont très peu lumineuses, et leur lumière se perd dans l'éblouissement de leur étoile hôte, telles des lucioles à côté d'un lampadaire. Ce *rapport de luminosité*, ou *contraste*, dépend de la longueur d'onde d'observation : dans l'infrarouge, où les jeunes planètes géantes brillent encore de leur propre processus de formation, ce rapport de luminosité est d'environ $10^{-5} - 10^{-6}$. Dans les longueurs d'onde visibles, où les planètes reflètent la lumière de leur étoile hôte, ce rapport devient encore plus difficile à surmonter puisqu'il peut atteindre 10^{-10} . En supposant qu'une planète doive ressembler étroitement à la Terre pour avoir une chance d'abriter la vie, un objectif important de l'imagerie directe est de trouver une planète rocheuse et relativement petite, orbitant autour d'une étoile de la séquence principale comme le Soleil à une distance d'une unité astronomique, qui est la distance entre le Soleil et la Terre. C'est à

³<https://exoplanets.nasa.gov/alien-worlds/ways-to-find-a-planet/#/2>

⁴<https://exoplanets.nasa.gov/alien-worlds/ways-to-find-a-planet/#/1>

partir de ces propriétés physiques de ces « *exoterras* » que nous déduisons les objectifs instrumentaux de contraste de 10^{-10} dans le visible à une séparation angulaire de 0,1 seconde d'arc.

Pour surmonter ces deux défis imposés par le comportement physique fondamental de la lumière, les astronomes utilisent des instruments complexes d'*imagerie à haut contraste* qui fournissent, comme leur nom l'indique, des capacités d'observation à haut contraste avec une haute résolution angulaire. Un tel instrument comporte trois composants principaux qui contribuent chacun à la performance globale de l'imagerie. Le premier, le *coronographe*, est un ensemble de masques optiques qui visent à réduire l'éblouissement global de l'étoile tout en préservant la lumière provenant de la planète. Le premier concept d'un tel coronographe a été construit et utilisé avec succès par l'astronome français Bernard Lyot en 1930. Il a utilisé une combinaison d'un *masque en plan focal* et de ce que nous appelons aujourd'hui un *diaphragme de Lyot* pour bloquer le disque du Soleil afin d'observer la faible couronne solaire. Ce n'est que plusieurs décennies plus tard que les scientifiques ont adapté ces instruments pour les diriger vers d'autres étoiles de manière à observer de faibles disques circumstellaires et des planètes.

La deuxième composante importante d'un instrument à haut contraste provient des distorsions auxquelles la lumière est exposée avant d'atteindre le détecteur. Alors que la lumière provenant d'un objet ponctuel distant se déplace sous la forme d'une onde plane, elle peut rencontrer en chemin des sources d'aberrations qui introduisent des *erreurs de front d'onde*. Cela peut être causé par l'atmosphère terrestre dans laquelle les gradients de température et de pression, ainsi que le vent, provoquent des turbulences importantes qui détériorent significativement la qualité de l'image finale. De plus, il existe des sources d'erreurs de front d'onde internes au télescope qui contribuent également à cet effet, comme les mauvais alignements des différents composants du système optique, ou les erreurs de polissage des surfaces optiques. Le dispositif qui atténue ces effets est un *système d'analyse et de contrôle du front d'onde*. Il se compose d'un analyseur de front d'onde qui détecte les aberrations introduisant les erreurs, et de composants de contrôle tels que des *miroirs déformables* qui appliquent des corrections en adaptant la forme de leur propre surface en fonction des mesures de l'analyseur. Avec le coronographe, ces boucles de contrôle sont utilisées pour moduler la lumière dans une partie désignée du plan image, créant ainsi une zone sombre de très haute contraste, le « *dark hole* ».

Le troisième grand élément de l'imagerie à haut contraste est le traitement de données. Bien qu'il soit préférable de séparer la lumière de la planète et de l'étoile avec un système optique approprié avant la détection, il existe des méthodes de réduction des données afin d'améliorer le contraste détecté dans l'image finale et ainsi repousser la limite de détection.

Ces technologies ont permis d'obtenir une abondance de résultats scientifiques passionnants, par exemple dans le domaine des disques circumstellaires (voir figure 1.2 du chapitre 1). La première détection d'exoplanète par imagerie directe a été réalisée en 2004, et plusieurs autres ont suivi. Les observations d'exoplanètes géantes et chaudes comme celles de la figure 1.3 sont utilisées pour contraindre leur orbite planétaire et en apprendre davantage sur leur atmosphère. Cependant, avec les instruments dont nous disposons aujourd'hui, nous ne sommes pas en mesure de voir les petites exoplanètes rocheuses qui pourraient potentiellement abriter la vie. Même avec les télescopes terrestres de classe 30 m qui seront mis en service au milieu des années 2030, également appelés « *extremely large telescopes* », ou ELTs, nous serons limités à des niveaux de contraste qui atteindront au mieux 10^{-9} , ce qui ne permettra pas de voir des exoplanètes semblables à la Terre autour d'étoiles de type solaire. Pour ouvrir la possibilité de détecter ces objets, nous devons envoyer nos télescopes dans l'espace afin d'éviter l'influence néfaste de l'atmosphère terrestre, et nous devons développer des instruments d'imagerie à haut contraste plus performants que ceux que nous utilisons aujourd'hui.

L'envoi dans l'espace d'un télescope capable d'imager des exoterras pose de nombreux défis. Pour atteindre une résolution angulaire et une sensibilité suffisante pour imager une faible source ponctuelle située à seulement 0,1 seconde d'arc d'une étoile brillante, l'observatoire doit inclure un miroir primaire de plusieurs mètres de diamètre, pour avoir une très grande surface de collecte de la lumière. La solution technique pour construire de grandes ouvertures de plus de 4 m dans l'espace consiste à les composer de segments hexagonaux individuels, assemblés pour former le miroir primaire total. Cela les rend plus légers que les équivalents monolithiques, plus faciles à fabriquer, et ils peuvent également être repliés pour se loger dans la fusée qui les emporte dans l'espace. Le premier télescope segmenté envoyé dans l'espace

sera le James Webb Space Telescope, avec 18 segments combinés pour donner un miroir d'un diamètre total de 6,5 mètres. Il nous fournira de précieuses opportunités d'apprentissage pour les futurs télescopes segmentés encore plus ambitieux. Alors qu'une grande ouverture de télescope nous donne la résolution angulaire et la sensibilité nécessaires, l'architecture globale du système d'imagerie à haut contraste définit les niveaux de contraste que nous pouvons atteindre pendant les observations.

Il existe toute une série de coronographes adaptés à l'imagerie directe depuis l'espace, mais chacun d'entre eux s'accompagne de compromis dans ses capacités respectives : différents angles de travail internes et externes (à quelle distance de l'étoile ils peuvent observer, anglais : « *inner and outer working angle* »), transmission hors axe (quelle quantité de lumière planétaire ils perdent en éliminant la lumière stellaire, anglais : « *planet throughput* ») ou compatibilité avec les segments de la pupille du télescope. Ils sont néanmoins tous conçus pour pouvoir atteindre un contraste de 10^{-10} dans des conditions parfaites - sans aberrations et sans évolution temporelle du système optique. Bien qu'il n'y ait pas d'atmosphère à évolution rapide dans l'espace qui introduise des aberrations, il existe encore plusieurs sources importantes d'erreurs de front d'onde pour un télescope spatial, auxquelles les coronographes sont très sensibles. Les imperfections dans la fabrication des composants et des surfaces optiques introduisent des distorsions qu'un système embarqué d'analyse et de contrôle du front d'onde peut corriger dans une certaine mesure. D'autres aberrations sont introduites par des déformations thermiques et mécaniques de la structure de l'observatoire qui dégradent le contraste. Ceci est particulièrement vrai pour les déformations provenant de la structure maintenant l'ensemble des segments du miroir primaire. Certaines de ces erreurs de front d'onde évoluent dynamiquement pendant les observations et la question se pose de savoir si nous pouvons maintenir un télescope segmenté aussi grand suffisamment stable pour que le contraste ne se dégrade pas au point de devenir aveugle aux exotéres. Le but de cette thèse est de développer des méthodes pour caractériser les limites nécessaires à ces variations d'erreur de front d'onde, en particulier pour les distorsions introduites par les segments du miroir primaire, et de préparer les futurs télescopes spatiaux pour une utilisation efficace des instruments à haut contraste dans l'espace.

Chapitre 2: La base théorique du tolérancement des erreurs de front d'onde pour les grands télescopes segmentés

Les désalignements entre les segments d'un grand télescope contaminent les images des observations coronagraphiques avec des tavelures (« *speckles* ») dans la zone sombre de l'image où nous voulons observer les exoplanètes (« *dark hole* »). Avant de pouvoir décider comment atténuer ce problème, il est important de caractériser l'effet exact des désalignements des segments sur le contraste coronagraphique. Le modèle PASTIS de propagation des aberrations liées aux désalignements entre segments à travers un coronographe convertit fidèlement ces erreurs de front d'onde en un indice de contraste moyen et fournit une base analytique qui nous permet d'inverser le problème. Au lieu de se demander : « Quel est le contraste sous une quantité x d'aberrations ? », nous nous demandons : « Quel niveau x d'erreurs de front d'onde pouvons-nous tolérer si nous avons besoin d'un contraste spécifique pour notre observation ? » Ce chapitre montre comment exprimer ces tolérances d'erreur de front d'onde avec des équations analytiques et présente leur interprétation statistique sur un ensemble de réalisations de désalignements du miroir segmenté. L'un des principaux résultats du travail effectué dans ce chapitre est que tous les segments d'un télescope segmenté équipé d'un coronographe ne doivent pas être contraints de manière aussi stricte.

Chapitre 3: Outils de simulation numériques et résultats simulés de tolérances

Les simulateurs numériques sont un outil indispensable pour l'étude et la mise en œuvre du tolérancement. J'ai développé un logiciel écrit en Python qui effectue l'analyse de tolérancement présentée au chapitre 2. Il contient une série de simulateurs optiques qui reproduisent la géométrie du télescope et du coronographe pour plusieurs modèles de télescopes, qui peuvent être connectés aux calculs de tolérancement des erreurs de front d'onde. Il permet également d'utiliser des simulateurs qui ne sont pas inclus dans ce logiciel pour effectuer l'analyse. L'application de ces méthodes théoriques et numériques au cas concret d'un coronographe sur le télescope spatial James Webb nous donne une idée de la manière dont les performances de contraste de cet instrument sont liées aux désalignements potentiels du miroir primaire du télescope. Plus près de nous, ce chapitre présente également des expériences de tolérancement simulées sur le banc HiCAT.

Chapitre 4: Validations expérimentales du tolérancement des erreurs de front d'onde segmentées sur HiCAT

Le banc HiCAT au STScI est une expérience de laboratoire optique destinée à la démonstration de technologies d'imagerie spatiale à haut contraste. Il contient un coronographe de Lyot, deux miroirs déformables continus pour la détection et le contrôle du front d'onde, et un miroir déformable segmenté qui reproduit les effets optiques d'un grand miroir primaire segmenté. Ses 37 segments peuvent être contrôlés individuellement en piston, en tip et en tilt pour introduire des désalignements. Ce chapitre présente les résultats d'expériences réalisées avec ce système optique, dans lequel les équations analytiques développées au chapitre 2 sont confirmées sur un instrument réel à haut contraste. Les conclusions de ce travail nous permettent de déterminer dans quelle mesure nous devons contrôler le miroir déformable segmenté si nous voulons atteindre des niveaux de contraste qui se situent au-delà des performances actuelles du banc HiCAT.

Chapitre 5: Maintenir les grands télescopes segmentés dans l'espace suffisamment stables pour l'imagerie exotéristre

Le cadre du tolérancement analytique développé au chapitre 2 n'est en fait pas limité aux effets de désalignement des miroirs segmentés. La méthode est également applicable aux erreurs de front d'onde qui s'étendent à toute la pupille du télescope en même temps, ou pour traiter des aberrations spécifiques qui proviennent de déformations thermiques et mécaniques, trouvées par modélisation. Cela ouvre toute une série d'applications pour ce traitement très fondamental des aberrations. Il s'agit notamment d'aller un peu plus loin, en examinant de plus près comment certaines aberrations du front d'onde évoluent dans le temps, et quelles seraient les mesures appropriées pour contrer l'effet de leur évolution sur le contraste coronographique. Ce chapitre présente d'abord quelques résultats quantitatifs du tolérancement des erreurs de front d'onde induites thermiquement, puis met en évidence le rôle de ces méthodes dans la conception globale des stratégies d'analyse et de contrôle du front d'onde des grands télescopes segmentés dans l'espace.

Annexe A: Démonstration de l'analyse et du contrôle du front d'onde sur un banc optique segmenté

Avant de commencer toute observation avec un télescope segmenté, il faut réduire les désalignements des segments au moyen de l'analyse et du contrôle du front d'onde. Pour les observations incluant un coronographe, les segments doivent être alignés à un niveau égal ou supérieur à celui requis par l'analyse des tolérances développées au chapitre 2. Le banc d'essai JOST segmenté, mais non coronographique, au STScI utilise un algorithme d'analyse et de contrôle de front d'onde appelé LAPD pour trouver la meilleure solution d'alignement de ce banc optique. Ce chapitre présente les résultats des expériences LAPD sur JOST sur un grand champ optique avec pour objectif particulier d'aligner le miroir segmenté du banc.

Annexe B: Coronographie sur un banc optique à ouverture segmentée en laboratoire

De nombreuses technologies de coronographie, d'analyse et de contrôle du front d'onde ont été développées afin d'atteindre un haut contraste sur les télescopes spatiaux. Leur fonctionnement est généralement meilleur dans le cas parfait d'ouvertures circulaires non obstruées par des miroirs secondaires, des araignées, et également dans des conditions parfaites lorsqu'il n'y a pas d'erreurs de front d'onde dans le système optique. Cependant, les observatoires vrais, qui ne sont pas parfaits, contiennent plusieurs caractéristiques qui rendent difficile l'obtention d'un DH à haut contraste, comme celles énumérées ci-dessus. L'objectif du banc optique HiCAT est d'étudier comment combiner les divers composants des instruments à haut contraste, tels que les senseurs, les contrôleurs et les techniques de suppression de la lumière des étoiles, dans un système fonctionnant efficacement sur des ouvertures comprenant des obstructions centrales et de la segmentation. Ce chapitre présente les principales motivations, le concept de ce banc optique, et explique les différentes configurations du banc qui ont été utilisées pour mener des expériences d'analyse et de contrôle du front d'onde avec un coronographe à ouverture segmentée.

La raison principale des efforts déployés dans l'instrumentation pour l'imagerie à haut contraste est d'améliorer les résultats scientifiques que nous pouvons obtenir avec nos télescopes. Pour en savoir plus sur l'existence, mais aussi sur l'occurrence globale d'exoplanètes semblables à la Terre dans notre voisinage galactique, nous cherchons à augmenter le *rendement en exoplanètes candidates*. Cela signifie que nous

devons trouver un nombre suffisamment important de planètes potentiellement habitables pour commencer à tirer des conclusions plus globales, par exemple pour déterminer si la Terre est une planète plutôt typique de notre galaxie ou non. Bien qu'il n'existe actuellement aucun télescope spatial qui utilise de coronographe à contrôle actif, les futures missions nous donneront cette capacité. Le télescope spatial monolithique Roman de 2,4 m sera le premier observatoire spatial à utiliser des miroirs déformables, ce qui constituera une étape intermédiaire. Les missions futures, comme le concept de télescope segmenté LUVOIR de 15 m, déploieront les moyens d'un observatoire d'imagerie d'exoterres.

My thanks go out to Raphaël Pourcelot for his dedicated help in translating this summary into French.

DEUTSCHE ZUSAMMENFASSUNG

Seit Jahrhunderten fragt sich die Menschheit, was jenseits unseres Heimatplaneten, der Erde, liegt. Schon antike Philosophen wollten die Frage beantworten, ob Planeten existieren könnten, die um andere Sterne kreisen als um unsere Sonne. Die ersten konkreten wissenschaftlichen Theorien zu dieser Problematik bildeten sich im 19. und dem frühen 20. Jahrhundert, aber die erste Entdeckung eines Planeten auf einer Umlaufbahn um einen Stern ähnlich zu dem unseren erfolgte im Jahre 1995. Dies zwang uns, diese extrasolaren Planeten, *Exoplaneten* genannt, in einem völlig neuen Licht zu sehen. Es stellte sich heraus, dass es nicht nur viel mehr Planetensysteme gibt als unser eigenes, sondern wir lernten auch, dass die Artenvielfalt dieser Planeten weitaus größer ist als jene, die wir von den Himmelskörpern in unserem eigenen Sonnensystem kennen. Bei so vielen fremden Welten in unserer galaktischen Nachbarschaft stellt sich die Frage: Gibt es da draußen auch anderes Leben?

Die Detektion von Exoplaneten benötigt ausgeklügelte Instrumente, die das Signal dieser schwer zu erkennenden Objekte einfangen: Dadurch, dass es so viel einfacher ist, den viel helleren Zentralstern zu beobachten als einen Planeten, waren die ersten produktiven Methoden, die tausende solcher neuen Objekte entdeckten, sogenannte *indirekte* Beobachtungsmethoden. Diese sammeln kein direktes Licht vom Planeten selbst, sondern messen Effekte, die es uns ermöglichen, die Existenz eines oder mehrerer Planeten in der Umlaufbahn um einen Stern abzuleiten. Die zwei ertragreichsten solcher Methoden sind die *Transitmethode*⁵ und die *Radialgeschwindigkeitsmethode*⁶. Wenn ein Exoplanet während eines Transits durch unsere Sichtlinie hindurch an seinem Zentralstern vorbei wandert, dann verringert das die von uns gemessene Lichtstärke des Sternes. Dieser Effekt ist messbar und wir können dabei den Radius des Planeten ermitteln. In der Radialgeschwindigkeitsmethode messen wir die periodische Verschiebung der Wellenlänge des Lichtes von dem Stern durch den Dopplereffekt. Diese Verschiebung ergibt sich dadurch, dass der Planet den Stern mit seinem gravitationellen Einfluss zu einer Taumelbewegung zwingt. Mit diesen Daten lässt sich eine ungefähre Masse des Planeten errechnen.

Tausende neue Exoplaneten wurden mit solchen indirekten Methoden detektiert. Um aber mehr als nur deren Radius und Masse ermitteln zu können, müssen wir jenes Licht analysieren, das diese Planeten selbst emittieren und reflektieren, anstelle des Lichtes ihrer jeweiligen Zentralsterne. Dies wird durch *direkte* Detektionsmethoden ermöglicht, die das überwältigende Licht eines Sternes blockieren, um Bilder des Planeten selbst aufzunehmen. Dadurch können wir eine Spektralanalyse von der Atmosphäre eines Planeten erstellen, um zu ermitteln, welche Moleküle darin enthalten sind und ob außerirdisches Leben auf dem Planeten möglich wäre. Diese Methode nennen wir *direkte Beobachtung* (engl.: "direct imaging"). Dabei müssen einige Hürden überwunden werden, um sie zur Detektion von Exoplaneten anwenden zu können. Das erste Problem ist die geringe Distanz zwischen einem Exoplaneten und seinem Stern. Potenziell bewohnbare Exoplaneten befinden sich aus der Ferne betrachtet sehr nahe an ihrem Stern, mit einer projektierten Entfernung am Himmel von einer Bogensekunde bis zu wenigen Millibogensekunden. Um den Planeten also als einzelne Punktquelle auflösen zu können, benötigen wir ein sehr gutes *Winkelauflösungsvermögen*, das wir nur mit Teleskopen erreichen können, die sehr große Primärspiegel haben. Das zweite Problem ist die schwache Leuchtkraft dieser Himmelskörper. Diese Exoplaneten sind sehr lichtschwach und werden von dem um ein Vielfaches helleren Stern überstrahlt, ganz so wie ein Glühwürmchen direkt neben einem Flutlicht. Dieses *Helligkeitsverhältnis*, oder *Kontrast*, hängt von der Wellenlänge ab in der wir unsere Beobachtungen durchführen: Im infraroten Bereich, wo junge, riesige Planeten noch von ihrem eigenen Entstehungsprozess glühen, ist dieses Helligkeitsverhältnis etwa 10^{-5} – 10^{-6} . Im sichtbaren Licht, in dem Planeten das Licht ihres Zentralsternes reflektieren anstelle selber

⁵<https://exoplanets.nasa.gov/alien-worlds/ways-to-find-a-planet/#/2>

⁶<https://exoplanets.nasa.gov/alien-worlds/ways-to-find-a-planet/#/1>

Strahlung zu verursachen, vernehmen wir einen astrophysikalischen Kontrast von bis zu 10^{-10} . Wenn wir annehmen, dass ein Exoplanet unserer Erde ähneln muss, um fremde Lebensformen unterstützen zu können, dann müssen wir nach relativ kleinen Gesteinsplaneten suchen, die auf einer Distanz von einer astronomischen Einheit um einen sonnenähnlichen Stern kreisen, was der Entfernung von unserer Erde zur Sonne entspricht. Eines der Hauptziele der direkten Beobachtung ist daher solche sogenannte "Exoerden" zu finden. Um das zu erreichen, brauchen wir Instrumente, die unter solchen physikalischen Parametern Beobachtungen durchführen können: nämlich mit reflektiertem Licht im sichtbaren Bereich des Lichtspektrums, mit einem Kontrast von 10^{-10} auf einer Winkelentfernung von 0,1 Bogensekunde.

Um diese zwei Hürden, die aus dem grundlegenden, physikalischen Verhalten des Lichtes folgen, zu überwinden, benutzen Astronomen komplexe Instrumente für die Methode des "high-contrast imaging". Diese erlauben es uns, Beobachtungen mit einem hohen Lichtkontrast auf sehr kleinen Winkelentfernungen durchzuführen. Diese Instrumente bestehen aus drei Hauptkomponenten, die alle zur ganzheitlichen Leistung der Abbildungen beitragen. Die erste Komponente ist der *Koronograf*, der aus mehreren optischen Masken besteht, die gemeinsam eine Ablendung des Sternlichtes bewirken, während das Planetenlicht erhalten bleibt. Das erste Konzept für so einen Koronografen wurde von dem französischen Astronomen Bernard Lyot erstellt und im Jahre 1930 das erste Mal erfolgreich benutzt. Er verwendete eine Kombination aus einer *Brennpunktblende* (engl.: "focal-plane mask") und was wir heute eine *Lyotblende* nennen (engl.: "Lyot stop"), um die gleißende Oberfläche der Sonne abzudecken, was ihm die Beobachtung der lichtschwachen Sonnenkorona ermöglichte. Jahrzehnte später adaptierten Wissenschaftler solche Instrumente, um sie auf andere Sterne anzuwenden. Dies erlaubt uns heutzutage, lichtschwache zirkumstellare Scheiben und Exoplaneten zu beobachten.

Die zweite Komponente in einem Hochkontrastinstrument ist dadurch motiviert, dass das Licht Störungen erfährt, bevor es auf den Detektor trifft. Unangetastetes Licht von einer weit entfernten Punktquelle breitet sich als eine flache Wellenfront aus, aber wenn es unterwegs auf Aberrationsquellen trifft, erfährt es so genannte *Wellenfrontstörungen*. Dies kann durch die Erdatmosphäre verursacht werden, in der Temperatur- und Druckunterschiede sowie Windturbulenzen erzeugt werden, die die Qualität des aufgenommenen Bildes negativ beeinflussen. Zusätzlich gibt es viele Quellen für Wellenfrontstörungen innerhalb des Teleskopes selbst, wie zum Beispiel feine Fehler in der Ausrichtung von optischen Elementen oder Polierungsdefekte auf optischen Oberflächen. Eine Komponente, die zu einem gewissen Teil gegen diese Einflüsse helfen kann, ist ein *Wellenfrontsensorik und -kontrollsystem*. Es besteht aus einem Wellenfrontsensor, der die Abbildungsfehler im Lichtbündel detektiert, und Kontrollkomponenten wie *deformierbare Spiegel*, die Korrekturen anwenden. Das tun sie indem ihre Oberflächen an die Resultate der Messungen des Wellenfrontsensors angepasst werden, um die detektierten Fehler zu kompensieren. Zusammen mit dem Koronografen werden diese Kontrollschleifen benutzt, um das Licht in einem bestimmten Bereich der Bildebene zu modulieren und den Kontrast zu verbessern. Dadurch entsteht ein Bildbereich mit sehr hohem Kontrast, die sogenannte "dark hole".

Die dritte Komponente in der Hochkontrast-Bildgebung ist die Nachbearbeitung. Es ist selbstverständlich bevorzugt, das Planetenlicht mit optischen Komponenten von dem des Sternes zu trennen bevor es den Detektor erreicht. Allerdings gibt es sehr wohl Datenreduktionsmethoden, die es uns erlauben, den Kontrast auch in der Nachbearbeitung zu verbessern, was das Detektionslimit noch erweitert.

Diese Technologien ermöglichten zahlreiche aufregende wissenschaftliche Ergebnisse, zum Beispiel die Bilder von zirkumstellaren Scheiben in Abb. 1.2 in Kap. 1. Die erste Detektion eines Exoplaneten mit direkter Bildgebung erfolgte im Jahre 2004 und eine Handvoll anderer folgte. Beobachtungen von warmen Gasriesen wie die in Abb. 1.3 werden benutzt, um deren Umlaufbahnen zu bestimmen und mehr über ihre Atmosphären zu erfahren. Allerdings ist es uns mit den Instrumenten, die uns heute zur Verfügung stehen, nicht möglich, kleine Gesteinsplaneten zu beobachten, auf denen möglicherweise Leben besteht. Selbst mit den erdgebundenen Teleskopen der 30 m Klasse, den Riesenteleskopen oder sogenannten "extremely large telescopes" (ELTs), die im kommenden Jahrzehnt in Betrieb genommen werden, werden wir im besten Fall einen Kontrast von 10^{-9} erreichen. Dies hindert uns daran, erdähnliche Exoplaneten zu beobachten, die auf einer Umlaufbahn um sonnenähnliche Sterne kreisen. Um einen Weg zur Beobachtung dieser Objekte zu finden, müssen wir Teleskope im Weltraum platzieren, um die störende Erdatmosphäre zu vermeiden. Außerdem werden wir bessere Hochkontrastinstrumente entwickeln müssen, verglichen mit dem, was wir

heute haben.

Ein Teleskop in den Weltraum zu senden, das Exoerden abbilden kann, birgt einige Herausforderungen. Dieses muss die nötige Winkelauflösung und Lichtempfindlichkeit erreichen, damit ein lichtschwacher Planet abgebildet werden kann, der nur 0.1 Bogensekunden entfernt von einem hellen Stern liegt. Dazu muss der Primärspiegel des benutzten Teleskopes mehrere Meter groß sein, was sowohl sein Auflösungs-, als auch sein Lichtsammelvermögen steigert. Die praktische Lösung, um solche Teleskope mit Spiegeldurchmessern von mehr als ~ 4 m herzustellen ist, mehrere individuelle, hexagonale Segmente zu einem großen Primärspiegel zu verbinden. Das macht ihn leichter als äquivalente monolithische Spiegel, er ist einfacher herzustellen und kann auch zusammengefaltet werden, um in einer Trägerrakete Platz zu finden. Das erste segmentierte Teleskop im Weltraum ist das James-Webb-Weltraumteleskop, in dem seine 18 Segmente zu einem Spiegeldurchmesser von 6,5 Metern zusammengesetzt sind. Dieses Observatorium wird uns eine wertvolle Gelegenheit bieten, unser Wissen für zukünftige, noch ambitioniertere Teleskope mit neuer Erfahrung zu erweitern. Während uns der Teleskopdurchmesser die nötige Auflösung und Lichtsensitivität gibt, ist der generelle Aufbau eines Hochkontrastinstrumentes dafür verantwortlich, welche Kontrastlevel wir während der Beobachtungen erreichen können.

Es existiert eine ganze Reihe von Koronografmodellen, die eine direkte Bildgebung vom Weltraum aus ermöglichen, aber jedes kommt mit seinen eigenen Vor- und Nachteilen: Unterschiedliche innere und äußere Arbeitswinkel (wie nahe am Stern oder weit entfernt von ihm wir mit hohem Kontrast beobachten können, engl.: "*inner and outer working angle*"), Transmission abseits der optischen Achse (wie viel Planetenlicht geht verloren, während das Sternenlicht eliminiert wird, engl.: "*planet throughput*"), oder seine Kompatibilität mit Segmentierungsspalten in der Teleskopapertur. Nichtsdestotrotz sind alle von ihnen imstande, einen Kontrast von 10^{-10} zu erreichen, wenn die Umstände perfekt sind – also ohne Aberrationen und ohne eine zeitliche Entwicklung des optischen Systems. Obwohl wir im Weltraum nicht mehr gegen die Atmosphäre arbeiten müssen, birgt ein Weltraumteleskop immer noch Ursachen für signifikante Wellenfrontstörungen, gegenüber denen der Koronograf sehr empfindlich ist. Abweichungen in der Herstellung von optischen Komponenten und Oberflächen verursachen Distorsionen, die ein integriertes Wellenfrontsensorik und -kontrollsystem zu einem gewissen Teil kompensieren kann. Weitere Aberrationen stammen von thermischen und mechanischen Deformationen der Gesamtstruktur des Teleskopes, die den Kontrast negativ beeinflussen. Dies trifft im Speziellen auf Deformationen in der Rückwand des Primärspiegels zu, die die einzelnen Segmente zusammenhält. Manche von diesen Wellenfrontstörungen entwickeln sich dynamisch über die Dauer von Beobachtungen. Dies wirft die Frage auf, ob wir imstande sind, solche großen, segmentierten Teleskope stabil genug zu halten, ohne dass der Kontrast so weit degradiert, dass wir blind gegenüber erdähnlichen Exoplaneten werden. Das Ziel dieser Dissertation ist es, Methoden zu entwickeln, um die zulässigen Limitationen für die Veränderung solcher Wellenfrontstörungen zu ermitteln. Dabei wird der Fokus auf jene Abbildungsfehler gelegt, die von einer Fehlausrichtung der Segmente des Primärspiegels herrühren. Diese Arbeit trägt dadurch zur Vorbereitung von zukünftigen Weltraumteleskopen bei, die auf effiziente Abbildungen mit hohem Kontrast spezialisiert sind.

Kapitel 2: Die theoretische Basis des Tolerierens von Wellenfrontstörungen für große, segmentierte Teleskope

Die Fehlausrichtungen zwischen den Segmenten eines großen Teleskopes kontaminieren die Bilder einer koronografischen Beobachtung mit Lichtflecken, sogenannten "*speckles*". Diese tauchen genau in jenen Bereichen des Bildes auf, in denen wir Exoplaneten finden wollen, in der "*dark hole*". Bevor wir uns die Frage stellen, wie dieses Problem gelöst werden kann, muss der exakte Einfluss der unterschiedlichen Fehlausrichtungen der Segmente auf den koronografischen Kontrast charakterisiert werden. Das PASTIS Modell der Lichtausbreitung von segmentierten Aberrationen durch einen Koronografen wandelt diese Wellenfrontstörungen wahrheitsgetreu in einen mittleren Kontrast um und bereitet dadurch die analytische Basis vor, die es uns erlaubt, das Problem umzukehren. Anstelle zu fragen: „Was ist der Kontrast, wenn eine Menge x an Aberrationen vorliegt?“, fragen wir nun: „Welches Level x an Wellenfrontstörungen können wir tolerieren, wenn wir einen spezifischen Kontrast in unseren Beobachtungen benötigen?“ Dieses Kapitel veranschaulicht wie diese Toleranzen mit analytischen Gleichungen abgeleitet werden können und stellt ihre statistische Interpretation über viele Zustände des segmentierten Spiegels hinaus vor. Ein bedeutendes Ergebnis aus diesem Kapitel ist, dass nicht alle Segmente auf einem segmentierten Teleskop

mit einem Koronografen gleichmäßig beschränkt werden müssen.

Kapitel 3: Numerische Simulationswerkzeuge und simulierte Tolerierungsergebnisse

Numerische Simulatoren sind unabdingbare Werkzeuge für die Demonstration einer Idee in einem sorgfältig implementierten Modell. Das PASTIS Tolerierungspackage ist eine in Python geschriebene Softwareeinheit, die die Tolerierungsanalyse aus Kap. 2 durchführt. Es enthält eine Reihe an optischen Simulatoren, die die Teleskopgeometrien und die Koronografen einiger bestehender Observatoriumskonzepte wiedergeben. Diese wiederum können mit der Kalkulation von Wellenfrontstörungstoleranzen verbunden werden. Das Package enthält auch die Möglichkeit, externe Simulatoren, die nicht in dieser Software enthalten sind, mit der Analyse zu verbinden. Die Anwendung dieser theoretischen und numerischen Methoden an den konkreten Fall eines Koronografen auf dem James-Webb-Weltraumteleskop gibt uns ein Gefühl dafür, wie der Kontrast in so einem Instrument von den potenziellen Ausrichtungsfehlern zwischen den Primärspiegelsegmenten abhängt. In einer etwas näher liegenden Anwendung präsentiert dieses Kapitel auch simulierte Tolerierungsexperimente für die optische Testvorrichtung namens HiCAT.

Kapitel 4: Experimentelle Validationen von segmentierten Wellenfrontstörungstoleranzen auf HiCAT

Die Testvorrichtung HiCAT am STScI ist ein optisches Laborexperiment zur Demonstration von Technologien für Hochkontrastbeobachtungen aus dem Weltraum. Es besteht aus einem Lyot-Koronografen, zwei deformierbaren Spiegeln mit durchgehender Oberfläche für Wellenfrontsensorik und -kontrolle sowie einem segmentierten deformierbaren Spiegel, der die optischen Effekte eines großen, segmentierten Primärspiegels nachbildet. Seine 37 Segmente sind einzeln in Piston, Tip und Tilt kontrollierbar und werden dazu benutzt, um dem System Ausrichtungsfehler zuzuführen, um im Folgenden ihre Effekte zu untersuchen. Dieses Kapitel präsentiert die Ergebnisse von Experimenten mit diesem Aufbau, mit denen die analytischen Gleichungen, die in Kap. 2 entwickelt wurden, auf einem realen Hochkontrastinstrument validiert werden. Die Schlussfolgerungen dieser Experimente lassen uns ermitteln, wie präzise wir den segmentierten deformierbaren Spiegel kontrollieren können müssen, um ein Kontrastlevel zu erreichen, das momentan jenseits der Leistungsfähigkeit von HiCAT liegt.

Kapitel 5: Große, segmentierte Weltraumteleskope stabil genug halten, um Exoerden zu beobachten

Der analytische Tolerierungsrahmen, der in Kap. 2 entwickelt wurde, ist nicht auf Ausrichtungseffekte von segmentierten Spiegeln beschränkt. Diese Methode kann genauso gut auf Wellenfrontstörungen angewendet werden, die sich über die gesamte Pupillenebene des Teleskopes erstrecken, sowie auf Aberrationen, die von thermischen und mechanischen Deformationen stammen. Dies eröffnet eine Reihe neuer Möglichkeiten, diese fundamentale Behandlung von Abbildungsfehlern anzuwenden. Es bedeutet auch, einen Schritt weiterzugehen und sich anzusehen, wie sich gewisse Wellenfrontaberrationen über die Zeit hinweg entwickeln und was angemessene Maßnahmen sind, um dem Effekt von sich zeitlich entwickelnden Aberrationen auf den koronografischen Kontrast entgegenzuwirken. Dieses Kapitel präsentiert zuerst einige quantitative Ergebnisse der Tolerierung von thermal induzierten Aberrationen. Danach zeigt es die Rolle solcher Methoden in dem gesamtheitlichen Design von Wellenfrontsensorik und -kontrollstrategien für große, segmentierte Weltraumteleskope auf.

Appendix A: Demonstration von Wellenfrontsensorik und -kontrolle auf einer segmentierten optischen Testvorrichtung

Bevor jegliche Beobachtung mit einem segmentierten Teleskop angefangen werden kann, müssen die Ausrichtungsfehler zwischen den einzelnen Segmenten mittels Wellenfrontsensorik und -kontrolle so gut es geht reduziert werden. Für Beobachtungen mit einem Koronografen müssen die Segmente zu einem gleichen oder besseren Level ausgereicht werden, wie es die Tolerierungsmethode in Kap. 2 vorgibt. Die segmentierte, aber nicht koronografische Testvorrichtung JOST am STScI benutzt einen Wellenfrontsensorik und -kontrollalgorithmus namens LAPD, um die beste Lösung hierfür zu finden. Dieses Kapitel präsentiert die Ergebnisse von LAPD Experimenten auf JOST auf einem weiten Bildfeld, mit dem spezifischen Ziel, die Segmente in dieser Vorrichtung gleichmäßig aneinander auszurichten.

Appendix B: Koronografie auf einer optischen Testvorrichtung mit segmentierter Spiegelöffnung

Es wurden viele Technologien im Bereich der Koronografie sowie Wellenfrontsensorik und -kontrolle entwickelt, die es uns erlauben, einen hohen Kontrast auf Weltraumteleskopen zu erreichen. Diese funktionieren im Normalfall am besten auf solchen Spiegelöffnungen, die nicht von Sekundärspiegeln und deren Stützen beschattet werden, sowie in perfekten Konditionen, wenn keine Wellenfrontstörungen im optischen System vorhanden sind. Allerdings besitzen reale Teleskope eine ganze Reihe an Merkmalen, die das Erreichen eines hohen Kontrastes erschweren, so die oben gelisteten Eigenschaften. Das Ziel der optischen Testvorrichtung namens HiCAT ist es, zu untersuchen, wie man die verschiedenen Komponenten eines Hochkontrastinstrumentes wie Sensoren, Kontrollelemente und lichtblendende Technologien in ein einzelnes System zusammen zusammenfügt, das effizient mit bedeckten und segmentierten Aperturen funktioniert. Dieses Kapitel führt die Motivation und das Design von HiCAT auf. Es erklärt weiters die unterschiedlichen Aufbaumodi, die benutzt wurden, um Experimente für Wellenfrontsensorik und -kontrolle auf einem Koronografen mit segmentierter Spiegelöffnung durchzuführen.

Der maßgebende Grund, warum so viel Aufwand in die Entwicklung von Instrumentation für “high-contrast imaging” gesteckt wird, sind die wissenschaftlichen Ergebnisse, die wir mit unseren Teleskopen erhalten. Um mehr über die Existenz, aber auch die Häufigkeit von erdähnlichen Exoplaneten in unserer galaktischen Nachbarschaft zu lernen, müssen wir den *Ertrag von Exoplanetenkandidaten* erhöhen. Das bedeutet, dass wir eine ausreichend hohe Zahl an möglicherweise bewohnbaren Planeten finden müssen, um globalere Schlüsse über ihre Population ziehen zu können, zum Beispiel ob die Erde ein eher typischer Planet in unserer Galaxie ist oder nicht. Obwohl heutzutage kein Weltraumteleskop existiert, das aktiv kontrollierte Koronografen benutzt, werden zukünftige Teleskope mit solchen Fähigkeiten ausgestattet sein. Das monolithische 2,4 m Roman-Weltraumteleskop wird das erste weltraumgebundene Teleskop sein, das die Benutzung von deformierbaren Spiegeln demonstriert. Dadurch stellt es einen wichtigen Meilenstein in der Vorbereitung zukünftiger Weltraumteleskope dar, die groß genug und leistungsstark genug sein werden, um Exoerden zu beobachten, wie zum Beispiel das 15 m große LUVUOIR-Teleskop.

HRVATSKI SAŽETAK

Ljudi su se stoljećima pitali i zamišljali što se nalazi izvan našeg matičnog planeta, Zemlje. Stari filozofi pitali su se postoje li planeti koji kruže oko zvijezda koje nisu naše Sunce. Konkretno znanstvene teorije koje se bave ovim pitanjima pojavile su se kroz 19. i u ranom 20. stoljeću. Prvo otkriće planeta koji kruži oko zvijezde slične Suncu 1995. godine potaknulo nas je da vidimo ove ekstrasolarne planete, odnosno *egzoplanete*, na potpuno novi način. Ne samo da postoji puno više planetarnih sustava od našeg, već je utvrđeno da je raznolikost planeta vrlo različita u usporedbi s onim što znamo o tijelima u našem Sunčevom sustavu. S toliko stranih svjetova u prostranstvu svemira, postavlja se pitanje: postoji li život izvan Zemlje?

Otkrivanje egzoplaneta zahtijeva sofisticirane instrumente koji hvataju signale tih objekata koje je teško otkriti. Budući da je lakše promatrati puno svjetliju matičnu zvijezdu, a ne izravno slabo vidljiv planet, prve metode otkrivanja koje su otkrile tisuće novih planetarnih tijela su one koje nazivamo *neizravnim* metodama otkrivanja. Te metode ne prikupljaju svjetlost planeta izravno, već mjere efekte koji nam omogućuju da zaključimo postoji li jedan ili više planeta u orbiti oko promatrane zvijezde. Dvije najplodonosnije metode u tom smislu jesu metoda *planetarnih tranzita*⁷ i metoda *radijalne brzine*⁸. Kada egzoplanet prolazi kroz izravnu liniju između Zemlje i svoje matične zvijezde, možemo zamijetiti zatamnjeno nastalo u svjetlu zvijezde, iz čega možemo odrediti samo postojanje egzoplaneta i njegov polumjer. Metodom radijalne brzine promatra se periodični pomak valne duljine svjetlosti koja do nas dopire od zvijezde, a što je uzrokovano njezinim gibanjem naprijed-natrag uslijed gravitacijske sile kojom planet djeluje na nju; ovo nam također govori o približnoj masi egzoplaneta.

Tisuće egzoplaneta otkriveni su takvim neizravnim metodama, no da bismo mogli karakterizirati ove udaljene svjetove ne samo njihovim radijusom i masom, moramo dobiti izravan pristup svjetlosti koju oni reflektiraju ili emitiraju, umjesto da koristimo zvijezdu domaćina kao posrednika. Ovo nam omogućuje metoda *izravnog* slikanja, koja blokira ogromnu svjetlost zvijezde i snima slike samih planeta. To nam dopušta da napravimo spektralnu analizu njihove atmosfere, govoreći nam koje molekule sadrže i bi li mogle podržati izvanzemaljski život. Ova metoda naziva se *izravno snimanje* (eng. "direct imaging"), a postoji nekoliko prepreka koje treba prevladati kako bi se mogla koristiti za otkrivanje egzoplaneta. Prvo, kada se promatraju iz daljine, ovi su planeti vrlo blizu matične zvijezde s projiciranim razmacima na nebu od jedne kutne sekunde do nekoliko kutnih milisekundi. Razlučivanje planeta kao pojedinačnog točkastog izvora svjetlosti stoga zahtijeva dobru *kutnu razlučivost* koju omogućuju teleskopi s vrlo velikim zrcalima. Drugo, egzoplaneti su vrlo slabog sjaja i njihova svjetlost se gubi u odsjaju njihove matične zvijezde, slično kao da pokušavate vidjeti krijesnicu tik uz jarku uličnu rasvjetu. Ovaj *omjer svjetline*, ili *kontrast*, ovisi o valnoj duljini na kojoj promatramo: u infracrvenom području, gdje mladi, divovski planeti još uvijek emitiraju svjetlost od svojih vlastitih procesa formiranja, ovaj omjer svjetline je oko 10^{-5} – 10^{-6} . U vidljivim valnim duljinama, gdje planeti reflektiraju svjetlost svoje matične zvijezde, ovaj omjer postaje još teže prevladati jer doseže čak 10^{-10} . Pod pretpostavkom da planet mora jako ličiti Zemlji da bi imao priliku podržati život, moramo potražiti stjenovite i relativno male planete koji kruže oko zvijezde glavnog niza poput Sunca na udaljenosti od jedne astronomske jedinice, koliko iznosi udaljenost između Sunca i Zemlje. Značajan cilj izravnog snimanja je stoga pronaći takve "egzozemlje". Instrumenti za opservacije planeta s takvim fizičkim svojstvima moraju moći provoditi detekcije u reflektiranom svijetlu vidljivih valnih duljina, s kontrastom od 10^{-10} pri kutnom razmaku od 0,1 kutne sekunde.

Kako bi prevladali ova dva izazova koja su nametnuta temeljnim, fizikalnim ponašanjem svjetlosti, astronomi koriste složene instrumente za *snimanje visokog kontrasta* koji, kao što im i ime sugerira, pružaju mogućnosti promatranja s visokim kontrastom pri maloj kutnoj razlučivosti. Postoje tri glavne kompo-

⁷<https://exoplanets.nasa.gov/alien-worlds/ways-to-find-a-planet/#/2>

⁸<https://exoplanets.nasa.gov/alien-worlds/ways-to-find-a-planet/#/1>

nente u takvom instrumentu od kojih svaka doprinosi ukupnoj izvedbi slike. Prvi od njih, *koronograf*, predstavlja skup optičkih maski koje imaju za cilj smanjiti ukupni odsjaj zvijezde, a da pritom očuvaju svjetlost koju emitira planet. Prvi koncept za takav koronograf uspješno je izradio i upotrijebio francuski astronom Bernard Lyot 1930. godine, kada je upotrijebio kombinaciju *žarišne pločice* (eng. "focal-plane mask") i onoga što danas zovemo "*Lyot stop*" (otvor koji smanjuje učinak ogiba oko rubova žarišne pločice) kako bi blokirao površinu Sunca kako bi mogli promatrati slabu solarnu koronu. Samo nekoliko desetljeća kasnije znanstvenici su prilagodili takve instrumente da budu usmjereni prema drugim zvijezdama, omogućujući nam da promatramo slabe cirkumstelarne diskove - i planete.

Druga važna komponenta instrumenta visokog kontrasta proizlazi iz izobličenja kojima je svjetlost izložena prije nego što dođe do detektora. Dok neometana svjetlost udaljenog, točkastog objekta putuje kao ravni val, na putu može naići na izvore aberacija koje unose *pogrešku valne fronte*. To se može dogoditi zbog Zemljine atmosfere, u kojoj razlike temperature i tlaka, kao i vjetar, uzrokuju turbulencije koje značajno pogoršavaju kvalitetu konačne slike. Štoviše, postoje izvori pogrešaka valne fronte unutar teleskopa koji također pridonose ovom učinku, poput neusklađenosti različitih komponenti u optičkom sustavu ili pogrešaka poliranja na optičkim površinama. Komponenta koja do određenog stupnja ublažava ove učinke je *sustav za očitavanje i kontrolu valne fronte*. Sastoji se od senzora valne fronte koji detektira aberacije koje unose pogreške i upravljačkih komponenti poput *deformabilnih zrcala* koja vrše korekcije prilagođavajući oblike vlastite površine prema mjerenju senzora. Zajedno s koronografom, takve se kontrolne petlje koriste za modulaciju svjetlosti u određenom dijelu ravnine slike, stvarajući tamno područje vrlo visokog kontrasta, odnosno "*dark hole*".

Treća velika komponenta visokokontrastne slike je naknadna obrada. Iako je poželjno odvojiti svjetlost planeta i zvijezde odgovarajućim optičkim sustavom prije nego što budu detektirani, postoje metode redukcije podataka koje nam omogućuju da poboljšamo mjereni kontrast na konačnoj slici, unaprijeđujući granicu otkrivanja još više.

S ovim tehnologijama došlo je do obilja uzbudljivih znanstvenih rezultata, primjerice iz slika cirkumstelarnih diskova poput onih prikazanih na slici 1.2 u Pogl. 1. Prvo otkrivanje egzoplaneta izravnim snimanjem obavljeno je 2004. godine, a nekoliko drugih je uslijedilo naknadno. Promatranja toplih, divovskih egzoplaneta poput onih prikazanih na slici 1.3 koriste se za preciznije određivanje njihovih planetarnih orbita i saznanja o njihovoj atmosferi. Međutim, s instrumentima koje danas imamo, nismo u mogućnosti vidjeti male, stjenovite egzoplanete na kojima bi se potencijalno mogao nalaziti život. Čak i sa zemaljskim teleskopima klase 30 m koji će biti operativni krajem 2020-ih i sredinom 2030-ih, koji se također nazivaju ekstremno veliki teleskopi ili ELT (eng. "extremely large telescopes"), bit ćemo ograničeni na razine kontrasta koje će u najboljem slučaju doseći 10^{-9} , sprečavajući nas da vidimo egzoplanete nalik Zemlji oko zvijezda nalik Suncu. Kako bismo omogućili promatranje takvih objekata, moramo postaviti naše teleskope u svemir kako bismo izbjegli štetan utjecaj Zemljine atmosfere; i moramo razviti moćnije instrumente za snimanje s visokim kontrastom od onoga što danas koristimo.

Slanje teleskopa sposobnog za snimanje egzozemlje u svemir povezano je sa nizom značajnih izazova. Kako bi se postigle kutna rezolucija i potrebna osjetljivost koje su dovoljne za snimanje slabog točkastog izvora svjetla koji je udaljen samo 0,1 kutnu sekundu od svijetle zvijezde, opservatorij treba primarno zrcalo promjera od nekoliko metara što daje vrlo veliku površinu prikupljanja svjetlosti. Inženjersko rješenje za izgradnju velikih teleskopskih otvora, većih od ~4 m u promjeru, je njihovo sastavljanje od pojedinačnih, šesterokutnih segmenata, složenih da tvore ukupno primarno zrcalo. To ih čini lakšim od monolitnih ekvivalenata, jednostavnijim za proizvodnju, a također se mogu sklopiti kako bi stali u oklop rakete koji ih dovodi u svemir. Prvi segmentirani teleskop poslan u svemir je svemirski teleskop James Webb, s 18 segmenata koji zajedno tvore zrcalo ukupnog promjera od 6,5 metara, a pružit će nam vrijedne prilike učenja za buduće, još ambicioznije segmentirane teleskope. Dok nam veliki otvor teleskopa daje potrebnu kutnu razlučivost i osjetljivost, cjelokupna arhitektura visokokontrastnog slikovnog sustava definira koje razine kontrasta možemo postići tijekom promatranja.

Postoji cijeli niz modela koronografa koji su prikladni za izravno snimanje, no svaki od njih dolazi s prednostima i manama u odnosu na svoje sposobnosti: različiti unutarnji i vanjski radni kutevi (koliko blizu ili daleko od zvijezde mogu opažati, eng. "*inner and outer working angle*"), planetarna propusnost (koliko svjetlosti planeta izgube eliminirajući svjetlost zvijezde, eng. "*planet throughput*") ili kompatibilnost

sa segmentalnim prazninama u teleskopskoj zjenici. Ipak, svaki od njih može postići kontrast kvalitete 10^{-10} u idealnim uvjetima – bez aberacija i temporalnih evolucija optičkog sustava. Iako u svemiru ne postoji brzorazvijajuća atmosfera koja uvodi aberacije, i dalje postoje značajni izvori grešaka valne fronte na svemirskim teleskopima na koje je koronograf vrlo osjetljiv. Nesavršenosti u izradi optičkih komponenti i njihovih površina stvaraju distorzije koje ugrađeni sustav za očitavanje i kontrolu valnih fronti može ispraviti u određenoj mjeri. Dodatno, iako u svemiru ne postoji brzorazvijajuća atmosfera koja stvara aberacije, i dalje postoje aberacije nastale termalnim i mehaničkim deformacijama na cjelokupnoj opservacijskoj strukturi koje smanjuju kvalitetu kontrasta. Ovo je naročito istinito za deformacije koje izviru iz stražnje ravnine strukture primarnog zrcala, koja drži zajedno niz individualnih segmenata. Neke od ovih grešaka valnih fronti dinamično se razvijaju tijekom opservacije te se postavlja pitanje je li moguće takav velik, segmentirani teleskop očuvati dostatno stabilnim kako kontrast ne bi izgubio kvalitetu u tolikoj mjeri da bismo bili slijepi na egzozemlje. Cilj ovog rada je razviti metode koje karakteriziraju nužne limite za ove varijacije valnih fronti, posebice distorzije nastale na primarnim zrcalnim segmentima, kao i pripremiti buduće svemirske teleskope za efikasnu upotrebu visokokontrastnih instrumenata u svemiru.

Poglavlje 2: Teoretska osnova tolerancije grešaka valnih fronti za velike segmentirane teleskope

Neusklađenost segmenata velikog teleskopa zagađuje slike koronografskih zapažanja mrljama svjetlosti (eng. “*speckles*”) na točno onim dijelovima slike na kojima bismo htjeli promatrati egzoplanete (“dark hole”). Prije nego li odlučimo kako doskočiti ovom problemu, važno je točno okarakterizirati efekt nedostatka poravnanja segmenata na koronografski kontrast. PASTIS model za propagaciju segmentiranih aberacija kroz koronograf vjerodostojno pretvara greške valnih fronti u brojni prikaz prosječne količine kontrasta te daje analitičku osnovu koja nam omogućava da preokrenemo problem. Pitanje „Koji je kontrast u slučaju x količine aberacija?” postaje „Koja x količina grešaka valnih fronti može biti tolerirana ako trebamo specifičan kontrast za svoje promatranje?” Ovo poglavlje pokazuje kako izvesti tolerancije grešaka valnih fronti analitičkim jednadžbama i uvodi njihovu statističku interpretaciju kroz niz različitih ostvarenih nedostataka poravnanja dijelova segmentiranog zrcala. Jedan od važnih rezultata proizašlih iz rada u ovom poglavlju jest da nije nužno sve segmente na segmentiranom zrcalu koje nosi koronograf jednako strogo ograničiti.

Poglavlje 3: Numerički simulacijski alati i simulirani rezultati tolerancije

Numerički simulatori nezamjenjiv su alat za demonstraciju ideje na pažljivo implementiranom modelu. PASTIS paket za toleriranje je softver napisan u Pythonu koji izvodi tolerancijsku analizu prezentiranu u Pogl. 2. Sadrži niz raznih optičkih simulatora koji reproduciraju geometriju i koronografe nekoliko različitih teleskopa koji mogu biti spojeni u kalkulacije tolerancije grešaka valnih fronti. Dodatno, on omogućava spajanje vanjskih simulatora koji nisu u PASTIS paketu u taj proces. Time se omogućava vršenje iste analize tolerancije grešaka valnih fronti za dodatne simulatore. Primjena ove teoretske i numeričke metode na konkretan slučaj koronografa na svemirskom teleskopu James Webb daje nam ideju kako njegova razina kontrasta ovisi o potencijalnim nepravilnostima u poravnanju primarnog zrcala ovog teleskopa. Naposljetku, ovo poglavlje pokazuje i simulirane tolerancijske eksperimente na optičkom ispitnom stolu HiCAT.

Poglavlje 4: Eksperimentalna potvrda tolerancije grešaka segmentiranih valnih fronti na HiCAT-u

HiCAT ispitni stol u STScI-u je optički laboratorijski eksperiment za demonstraciju tehnologija visokokontrastnog mapiranja iz svemira. Sadrži Lyot koronograf, dva neprekidna deformabilna zrcala za očitavanje i kontrolu valnih fronti, te segmentirano deformabilno zrcalo koje reproducira optička svojstva velikog segmentiranog primarnog zrcala. Njegovih 37 segmenata mogu se individualno kontrolirati na visinu, i nagibno na dvije međusobno okomite osi. Time se stvaraju nedostaci u poravnanju koje onda izučavamo. Ovo poglavlje prikazuje rezultate eksperimenata s ovim parametrima u kojima se analitičke jednadžbe iz Pogl. 2 potvrđuju na pravom visokokontrastnom instrumentu. Zaključci iz ovog rada omogućavaju nam da utvrdimo koliko dobro moramo moći kontrolirati segmentirana deformabilna zrcala ako želimo postići razine kontrasta koje su bolje od trenutnih mogućnosti HiCAT-a.

Poglavlje 5: Kako održati velike segmentirane teleskope u svemiru dovoljno stabilnima za direktno promatranje egzozemalja

Okvir analitičke tolerancije razvijene u Pogl. 2 nije ograničena samo na učinak nedostatka poravnjanja u segmentiranim zrcalima. Metoda je jednako primjenjiva na greške valnih fronti nastalih preko cijele zjenice teleskopa odjednom, ili za ispravljanje konkretnih aberacija koje nastaju uslijed termalnih i mehaničkih deformacija, a koje pronalazimo modeliranjem. Ovo otvara cijelu paletu mogućnosti primjene za ovaj fundamentalni način obrade aberacija. To uključuje i odlazak korak unaprijed, gledajući pobliže na to kako pojedine aberacije valnih fronti evoluiraju kroz vrijeme te koji bi bio prikladan postupak za otklanjanje učinka promjenjivih aberacija na koronografski kontrast. Ovo poglavlje prvo prezentira kvantitativne rezultate toleriranja termalno induciranih grešaka valnih fronti te zatim pokazuje ulogu takvih metoda u cjelokupnom dizajnu strategija za očitavanje i kontrolu valnih fronti na velikim segmentiranim teleskopima u svemiru.

Prilog A: Demonstracija očitavanja i kontrole valnih fronti na segmentiranom ispitnom stolu

Prije nego li bilo kakvo promatranje segmentiranim teleskopom može biti započeto, nedostaci u poravnanju moraju biti reducirani pomoću očitavanja i kontroliranja valnih fronti. Za promatranje koje uključuju koronograf, segmenti moraju biti poravnati u mjeri koja je jednaka ili bolja od zahtjeva postavljenih tolerancijskom analizom u Pogl. 2. Segmentirani, ali ne koronografski ispitni stol JOST u STScI-u koristi algoritam za očitavanje i kontrolu valnih fronti zvan LAPD kako bi pronašao optimalno rješenje poravnjanja ove optičke osnove. Ovo poglavlje prezentira rezultate LAPD eksperimenta na JOST-u na širokom vidnom polju s ciljem da poravna segmentirano zrcalo ovog ispitnog stola.

Prilog B: Koronografija na laboratorijskom ispitnom stolu sa segmentiranim zrcalom

Razvijeno je mnogo tehnologija za koronografiju te očitavanja i kontrole valnih fronti kako bi se postigao visoki kontrast na svemirskim teleskopima. One uglavnom rade najbolje u idealnom slučaju okruglog, monolitnog teleskopskog ogledala bez zatamnjenja uzrokovanog sekundarnim zrcalom i njegovim podupiračima te kada nema aberacija u optičkom sistemu. Međutim, stvarni teleskopi sadrže razna svojstva koja otežavaju postizanje visokog kontrasta, kao što su ova ranije navedena. Cilj HiCAT ispitnog stola je istraživanje u svrhu spajanja raznih sastavnica visokokontrastnog instrumenta kao što su senzori, kontroleri i tehnike suzbijanja svjetla kako bi sastavili učinkovit cjelokupni sustav koji uspješno radi i na zaklonjenim zjenicama sa segmentacijom. Ovo poglavlje navodi glavnu motivaciju i konstrukciju ovog ispitnog stola i objašnjava razne modove koji su korišteni za pokuse očitavanja i kontrole valnih fronti na koronografu sa segmentiranim otvorom.

Krajnji razlog za trud uperen prema instrumentaciji visokokontrastnih promatranja je poboljšanje znanstvenih rezultata koje možemo ostvariti svojim teleskopima. Kako bismo saznali za njihovo postojanje, tj. samu učestalost pojavljivanja zemljolikih egzoplaneta u svom galaktičkom susjedstvu, ono što želimo povećati je učestalost opažanja kandidata egzoplaneta. To znači da želimo pronaći dovoljno velik broj potencijalno nastanjivih planeta kako bismo mogli izvesti globalne zaključke, poput toga je li Zemlja u biti tipičan planet u našoj galaksiji ili nije. Iako trenutno ne postoji svemirski teleskop koji rabi aktivno kontrolirane koronografe, buduće misije stvorit će nam tu mogućnost. Monolitski Roman Space Telescope od 2,4 m bit će prvi opservatorij u svemiru koji će koristiti deformabilna zrcala te kao takav predstavlja važan korak u daljnjem razvoju budućih misija, kao što je segmentirani 15-metarski LUVOIR teleskop čijim će se lansiranjem postaviti osnove za direktno promatranje egzozemalja.

I am truly grateful to Dorotea Sremec and Fran Alfirević for translating this summary into Croatian.

PUBLICATIONS AND ACTIVITIES BEYOND RESEARCH

Refereed first-author publications:

3. *Wavefront tolerances of space-based segmented telescopes at very high contrast: Experimental validation*
Laginja, I., Sauvage, J.-F., Mugnier, L.M., Pueyo, L., Perrin, M.D., Noss, J., Will, S.D., Brooks, K.J., Por, E.H., Petrone, P., Soummer, R.
2022, *Astronomy & Astrophysics*, 658, A84
2. *Analytical tolerancing of segmented telescope co-phasing for exo-Earth high-contrast imaging*
Laginja, I., Soummer, R., Mugnier, L.M., Pueyo, L., Sauvage, J.-F., Leboulleux, L., Coyle, L., Knight, J. S.
2021, *Journal of Astronomical Telescopes, Instruments, and Systems* 7(1), 015004
1. *ExoTiC-ISM: A Python package for marginalised exoplanet transit parameters across a grid of systematic instrument models*
Laginja, I., Wakeford, H. R.
2020, *Journal of Open Source Software*, 5(51), 2281

Other refereed publications:

6. *LUVOIR A segment-level thermal stability analysis*
Sahoo, A., Pueyo, L., **Laginja, I.**, Soummer, R., Coyle, L., Knight, J.S., East, M.
In preparation
5. *High-order coronagraphic wavefront control with algorithmic differentiation: first experimental demonstration*
Will, S.D., Perrin, M.D., Por, E.H., Noss, J., Sahoo, A., Petrone, P., **Laginja, I.**, Pourcelot, R., Redmond, S.F., Pueyo, L., Groff, T.D., Fienup, J.R., Soummer, R.
In preparation
4. *Implementation of a dark zone maintenance algorithm for speckle drift correction in a high contrast space coronagraph*
Redmond, S.F., Pogorelyuk, L., Pueyo, L., Por, E.H., Noss, J., Will, S.D., **Laginja, I.**, Brooks, K.J., Maclay, M., Fowler, J., Kasdin, N.J., Perrin, M.D., Soummer, R.
Submitted to JATIS
3. *Low-order wavefront control using a Zernike sensor through Lyot coronagraphs for exoplanet imaging: Blind stabilization of an image dark hole*
Pourcelot, R., N'Diaye, M., Soummer, R., Por, E.H., **Laginja, I.**, Carbillet, M., Benard, H., Brady, G., Canas, L., Dohlen, K., Lai, O., Noss, J., Perrin, M.D., Petrone, P., Pueyo, L., Redmond, S.F., Sahoo, A., Vigan, A., Will, S.D.
Submitted to *Astronomy & Astrophysics*
2. *The β Pictoris b Hill sphere transit campaign. Paper I: Photometric limits to dust and rings*
Kenworthy, M.A., Mellon, S.N., Bailey, J.I. III, Stuik, R., Dorval, P., Talens, G.J.J., Crawford, S.M.,

Mamajek, E.E., **Laginja, I.**, Ireland, M., Lomberg, B.B.D., Kuhn, R.B.
2021, *Astronomy & Astrophysics*, 648, A15

1. *bRing: An observatory dedicated to monitoring the β Pictoris b Hill sphere transit*
Stuik, R., Bailey, J.E., Lomberg, B.B.D., Dorval, P., Talens, G.J.J., Mellon, S.N., **Laginja, I.**, Rieder, S., Crawford, S.M., Mamajek, E.E., Kenworthy M.A.
2017, *Astronomy & Astrophysics*, 607, A45

First-author SPIE papers, conference proceedings:

3. *Predicting contrast sensitivity to segmented aperture misalignment modes for the HiCAT testbed*
Laginja, I., Soummer, R., Mugnier, L.M., Pueyo, L., Sauvage, J.-F., Leboulleux, L., Coyle, L., Knight, J.S., Perrin, M.D., Will, S.D., Noss, J., Brooks, K.J., Fowler, J.
2020, Proc. SPIE 11443, *Space Telescopes and Instrumentation 2020: Optical, Infrared, and Millimeter Wave*; 11443J
2. *Wavefront error tolerancing for direct imaging of exo-Earths with a large segmented telescope in space*
Laginja, I., Leboulleux, L., Pueyo, L., Soummer, R., Sauvage, J.-F., Mugnier, L., Coyle, L., Knight, J.S., St-Laurent, K., Por, E., Noss, J.
2019, Proc. SPIE 11117, *Techniques and Instrumentation for Detection of Exoplanets IX*, 1111717
1. *James Webb Space Telescope Optical Simulation Testbed V: Wide-field phase retrieval assessment*
Laginja, I., Brady, G., Soummer, R., Egron, S., Lajoie, C.-P., Bonnefois, A., Michau, V., Choquet, É., Ferrari, M., Leboulleux, L., Levecq, O., N'Diaye, M., Perrin, M.D., Petrone, P., Pueyo, L., Sivaramakrishnan, A.
2018, Proc. SPIE 10698, *Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave*, 106983N

Other SPIE papers, conference proceedings:

11. *Experimental validation of active control of low-order aberrations with a Zernike sensor through a Lyot coronagraph*
Pourcelot, R., N'Diaye, M., Por, E.H., Perrin, M., Soummer, R., **Laginja, I.**, Sahoo, A., Carbillet, M., Brady, G.R., Dohlen, K., Maclay, M., McChesney, E., Noss, J., Petrone, P., Pueyo, L., Vigan, A., Will, S.D.
2021, Proc. SPIE 11823, *Techniques and Instrumentation for Detection of Exoplanets X*, 118231M
10. *Dark zone maintenance results for segmented aperture wavefront error drift in a high contrast space coronagraph*
Redmond, S.F., Pueyo, L., Pogorelyuk, L., Por, E., Noss, J., **Laginja, I.**, Brooks, K., Perrin, M.D., Soummer, R., Kasdin, N.J.
2021, Proc. SPIE 11823, *Techniques and Instrumentation for Detection of Exoplanets X*, 118231K
9. *Implementation of a broadband focal plane estimator for high-contrast dark zones*
Redmond, S.F., Pueyo, L., Pogorelyuk, L., Noss, J., Will, S.D., **Laginja, I.**, Kasdin, N.J., Perrin, M.D., Soummer, R.
2021, Proc. SPIE 11823, *Techniques and Instrumentation for Detection of Exoplanets X*, 118231Q
8. *Wavefront control with algorithmic differentiation on the HiCAT testbed*
Will, S.D., Perrin, M.D., Por, E.H., Noss, J., Sahoo, A., Petrone, P., **Laginja, I.**, Pourcelot, R., Redmond, S.M., Pueyo, L., Groff, T.D., Fienup, J.R., Soummer, R.
2021, Proc. SPIE 11823, *Techniques and Instrumentation for Detection of Exoplanets X*, 118230V

7. *Implementation of a dark hole maintenance algorithm for speckle drift in a high contrast space coronagraph*
Redmond, S.M., Kasdin, N. J., Pogorelyuk, L., Soummer, R., Pueyo, L., Perrin, M.D., Maclay, M., Noss, J., **Laginja, I.**, Will, S.D., Fowler, J.
2020, Proc. SPIE 11443, Space Telescopes and Instrumentation 2020: Optical, Infrared, and Millimeter Wave; 114432K
6. *Estimating low-order aberrations through a Lyot coronagraph with a Zernike wavefront sensor for exoplanet imaging*
Pourcelot, R., N'Diaye, M., Brady, G., Carillet, M., Dohlen, K., Fowler, J., **Laginja, I.**, Maclay, M., Noss, J., Perrin, M., Petrone, P., Por, E., Sauvage, J.-F., Soummer, R., Vigan, A., Will, S.
2020, Proc. SPIE 11443, Space Telescopes and Instrumentation 2020: Optical, Infrared, and Millimeter Wave; 1144346
5. *First error budget of a deployable CubeSat telescope*
Sauvage, J.-F., Schwartz, N., Vievard, S., Bonnefois, A., Velluet, M.-T., Correia, C., Cassaing, F., Fusco, T., Michau, V., Krapez, J.-C., Ferrari, M., **Laginja, I.**
2020, Proc. SPIE 11443, Space Telescopes and Instrumentation 2020: Optical, Infrared, and Millimeter Wave; 1144330
4. *Phase-retrieval-based wavefront metrology for high-contrast coronagraphy: 2. Reconstructions through a shaped pupil apodizer*
Brady, G.R., Petrone, P., **Laginja, I.**, Brooks, K., Zhang, M., N'Diaye, M., Moriarty, C., Hagopian, J., Soummer, R.
2019, Proc. SPIE 11117, Techniques and Instrumentation for Detection of Exoplanets IX, 1111712
3. *The LUVUOIR Extreme Coronagraph for Living Planetary Systems (ECLIPS) I: searching and characterizing exoplanetary gems*
Pueyo, L., Stark, C., Juanola-Parramon, R., Zimmerman, N., Bolcar, M., Roberge, A., Arney, G., Ruane, G., Riggs, A.J., Belikov, R., Sirbu, D., Redding, D., Soummer, R., **Laginja, I.**, Will, S.
2019, Proc. SPIE 11117, Techniques and Instrumentation for Detection of Exoplanets IX, 1111703
2. *High-contrast imager for complex aperture telescopes (HiCAT): 5. first results with segmented-aperture coronagraph and wavefront control*
Soummer, R., Brady, G.R., Brooks, K., Comeau, T., Choquet, É., Dillon, T., Egron, S., Gontrum, R., Hagopian, J., **Laginja, I.**, Leboulleux, L., Perrin, M.D., Petrone, P., Pueyo, L., Mazoyer, J., N'Diaye, M., Riggs, A.J.E., Shiri, R., Sivaramakrishnan, A., St.Laurent, K., Valenzuela, A.-M., Zimmerman, N.T.
2018, Proc. SPIE 10698, Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave, 106981O
1. *Phase-retrieval-based wavefront metrology for high contrast coronagraphy*
Brady, G.R. , Moriarty, C., Petrone, P., **Laginja, I.**, Brooks, K., Comeau, T., Leboulleux, L., Soummer, R.
2018, Proc. SPIE 10698, Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave, 106986I

Oral presentations

Git for Astronomers Intro

Self-built half-day introduction course on version control

17 March 2021; 14 April 2021

Virtual

Advanced Git for Astronomers

Self-built half-day course on advanced version control

5 May 2021; 19 May 2021

Virtual

Laboratory Demonstration of High Contrast Imaging on Segmented Apertures: Results from the STScI HiCAT Testbed

In the Spirit of Lyot, October 2019

Tokyo, Japan

Wavefront error tolerancing for direct imaging of exo-Earths with a large segmented telescope in space

SPIE Optics & Photonics, August 2019

San Diego, California, USA

Contrast-driven segment tolerancing for the co-phasing of segmented high-contrast instruments

NYRIA Workshop, October 2020

Virtual

Journées des doctorants de l'ONERA

2019, 2020, 2021

LAM PhD day

2021

Seminar talks

GRD (Marseille) 2018, LESIA (Paris) 2020, MPO (Nice) 2021

Poster presentations

Predicting contrast sensitivity to segmented aperture misalignment modes for the HiCAT testbed

SPIE Astronomical Telescopes & Instrumentation, December 2020

Virtual

Active Optics performance for the LUVOIR high contrast imager – Modelization and WFS experimental results

Adaptive Optics for Extremely Large Telescopes 6, June 2019

Quebec, Canada

Extending multi-aperture geometric alignment with ELASTICS to an 18 sub-aperture system

AAS meeting #233, January 2019

Seattle, Washington, USA

James Webb Space Telescope optical simulation testbed V: wide-field phase retrieval assessment

SPIE Astronomical Telescopes & Instrumentation, June 2018

Austin, Texas, USA

James Webb Space Telescope Optical Simulation Testbed: Segmented Mirror Phase Retrieval Testing

AAS meeting #231, January 2018

National Harbor, Maryland, USA

Software

During my PhD, I wrote a significant amount of code and contributed to various projects, many of which are open-source. The below table lists GitHub repositories I have made contributions to, roughly ordered by amount of contributed code. The PASTIS package contains the bulk of my PhD work.

Name	GitHub	
PASTIS	spacetelescope/PASTIS	Analytical tolerancing of coronagraphs
hcat-package	spacetelescope/hcat-package	Control code and simulator for HiCAT (private)
jost-package	spacetelescope/jost-package	Control code for JOST (private)
ExoTiC-ISM	hrwakeford/ExoTiC-ISM	Marginalized exoplanet transit parameters
FouFourier	ivalaginja/FouFourier	Learning notebooks for Fourier optics
catkit	spacetelescope/catkit	Hardware control interfaces for lab instrumentation
hccipy	ehpor/hccipy	Optical propagations for HCI
poppy	spacetelescope/poppy	Physical optical propagation

I also became an advocate for standardized software development principles in the astronomy community and designed my own course for teaching version control with git and GitHub. Split in an introductory and an advanced session, the first installments of this course took place in March and April of 2021 as a remote class with interactive tutorials. I want to thank Christopher Moriarty and Pablo Rodriguez Robles for their inputs and ideas for this course and their hands-on support during the classes.

Mentoring of students and interns

I had the pleasure to mentor and (co-)supervise several students for their internships during my PhD, both while at STScI as well as during my time at LAM.

2021	David Bourgeois	Master OSAE Paris, France	6 months/LAM
2020	Kelsey Glazer	Towson University, Baltimore, USA	6 months/STScI
2019	Lucas Batista	University of São Paulo, Brazil	2 months/STScI
2019	Maggie Kautz	University of Arizona, Tucson, USA	3 months/STScI

Community engagement

During the academic year 2020/2021, I was in charge of the organization of the weekly GRD group seminar together with my colleague Felipe Pedreros. I invited researchers from the international community to give presentations, coordinated schedules and hosted talks.

In early 2020, I was invited to serve a three year term on the **Equity, Diversity and Inclusion committee of SPIE**, the international society for optics and photonics. As part of this work, I have been involved in the development and discussion of strategies to make the global optics community a more welcoming and inclusive professional collective.

CURRICULUM VITAE

I was born on 17th of October 1992 in the city of Zagreb, Croatia. With my parents, Ana and Igor, and my younger sister, Marta, I grew up in Vienna, Austria. The frequent travels between Austria and Croatia for family visits and vacations incited an early interest in comparing the German and Croatian languages, both of which I grew up to use as a native speaker, and I started learning English as a foreign language at age five. I spent my summers on the beautiful Croatian coast with my aunt Ivana and my grandparents: Drago and Rajna on my mother's side, and Danka on my father's side. As a preschooler, it was there that I learned how to count: both by running rounds around our house, meticulously keeping track of the laps I did, clearly showing my early enthusiasm for sports. And also, by counting the stars appearing in the evening sky before I had to go to bed. I kept describing this activity with a childish skew of the Croatian word for stars, "zvijezde", fabricating a phrase that my late grandmother would lovingly repeat in later years: "Gledamo žveždice!" – "We are watching the stars!"

Between the four years in elementary school, and the eight years in high school at the Theresianische Akademie in Vienna, my interest would latch on to any new topic I laid my eyes on. My parents provided me with the opportunities I needed to explore them all, from language courses, over music classes and summer schools to athletic training and competition. I was lucky to have had a range of enthusiastic and caring teachers, above all my high school head teacher Mag. Gábor Magyary-Kossa. He not only taught me French, but also the importance of critical thinking, the value of imagination and the reward from approaching life with serenity.

With my first big passion being competitive sports and after some stints in football, basketball, and field hockey, I embraced the sport of table tennis. For a decade between ages 13-23, my coaches Franz Klaus and Ivana Masařikova helped me not only to gain the technical finesse of this sport, but also to strengthen the self-confidence I carried into my adult life. My second main passion was an insatiable interest in technical topics, specifically the natural sciences. Astronomy in particular sparked a hunger for knowledge that I could never fully satisfy with the books I read during my time in high school, so after graduating in 2011 I enrolled in the University of Vienna to pursue a bachelor's degree in astronomy. To solidify my knowledge of fundamental physics, I also enrolled in the physics bachelor's degree shortly after.

A random glance at a poster at the Institute for Astrophysics in Vienna turned out to be a career-defining moment for me. It was the announcement for the 2014 Dunlap Institute Summer School on Astronomical Instrumentation at the University of Toronto, happening in August that year. During this week-long experience I met enthusiastic lecturers and inspiring students from all continents, all sharing a passion for the technical aspect of astronomy. Consequently, upon my return to Vienna I started my search to find where I could enter the world of instrumentation after my bachelor's degree. Between becoming the Vienna table tennis women's single champion in 2014 and an "Erasmus" exchange semester in Strasbourg, France, in the spring of 2015, I found that place to be Leiden University in the Netherlands, where I applied for and got accepted to the master's degree in astronomy and instrumentation.

I enrolled in the master's degree in the fall of 2015, after graduating with a BSc in astronomy that summer; I also finished the BSc in physics, remotely, in early 2016. I spent the following two years with a fantastic cohort of both international and Dutch students, and my time in Leiden would prove to be the base on which I built my knowledge and career later on. In parallel to our courses, in my first year I worked with Dr. Matthew Kenworthy on the modeling of exomoons, moons around exoplanets, and their orbits within exoplanetary disks. I was looking at what potential exomoon signatures could look like in the Hill sphere transit of β Pictoris b , which I wrote up as my minor thesis. In my second year, I prepared my master's thesis with Prof. Dr. Christoph Keller and Dr. Michael Wilby, working on simulations and laboratory demonstrations of the vector Apodizing Phase Plate (vAPP) coronagraph. This was my first

concrete exposure to coronagraphy and laboratory work, and I defended successfully in August of 2017.


While I considered PhD positions at this point, my gut feeling was telling me that I was better off following a different path. It was only years later that I could put concrete words as to why and how that feeling manifested, but this decision brought me to the Space Telescope Science Institute (STScI) in Baltimore, USA, where I joined the Russell B. Makidon Optics Laboratory under Dr. Rémi Soummer as a Research and Instrument Analyst, later promoted to Astronomical Optics Scientist. It was an exciting time to join the institute, with the intended launch of the James Webb Space Telescope (JWST) only a year away, planned for the fall of 2018 at the time. I spent my time working on wide-field wavefront sensing and control on the JOST testbed, until about six months after joining STScI my manager and I started talking about my longer-term career options. Considering how my enthusiasm and joy with the work I had been doing at the lab was only growing by the day, I decided to buy into my then-manager's idea to create a suitable PhD position. We identified a suitable collaboration with Dr. Laurent Mugnier from the Office national d'études et de recherches aérospatiales (ONERA) in Châtillon, near Paris, and Dr. Jean-François Sauvage at the Laboratoire d'Astrophysique de Marseille (LAM) in France to craft a PhD that was co-funded by ONERA and STScI.

I remained in Baltimore for the first half of my 3-year doctorate which officially started on 1st October 2018 – shortly after crowning myself the 2018 Maryland State Women's Singles Champion in table tennis. I removed myself from JWST activities, whose launch was now scheduled for some time in 2021, around the same time as I was supposed to graduate. It turned into a running gag that I started working with the Makidon lab team just months before its launch, and that I would leave the group just before launch too - I set up a personal bet with myself to see who would do their thing first, me finishing my PhD or JWST being shot into space. The final result ended up way tighter than anticipated.

My main PhD research activities revolved around coronagraphy with segmented aperture telescopes, and I analyzed, simulated, and tested the impact of segment-level WFE on the coronagraphic contrast. With initial guidance from Dr. Lucie Leboulleux, who originally developed the PASTIS model for contrast calculations from aberrated segments, and in collaboration with Dr. Laurent Pueyo at STScI, and Dr. Laura Coyle and Dr. Scott Knight from Ball Aerospace, I developed a way to determine segment-level WFE tolerances for a given target contrast, and I provided this model with a statistical context. At the same time, as a member of the Makidon lab team, I started contributing heavily to the operations of the High-contrast imager for Complex Aperture Telescopes (HiCAT) testbed, which I used in the later stages of my PhD to perform experiments for my WFE tolerancing model. During my studies, I took part in various summer schools and conferences in places around the world, from Seattle in Washington/USA, San Diego in California/USA over Austin in Texas/USA and Quebec, Canada, to so many more. I was also trusted to give the Makidon lab overview presentation at the Spirit of Lyot conference in Tokyo, Japan, in October of 2019.

Many of these activities came to a sudden halt with the outbreak of the Coronavirus pandemic in March of 2020, just two weeks before I was set to move to Marseille for the second half of my PhD. With no flights to bring me to Europe, I stayed in Baltimore for a couple more weeks, during which I acted as the de facto deputy team lead to the Makidon lab while we were all trying to cope with the new situation. In May I got back to Europe, albeit to Austria since France had closed its borders to limit the circulation of the virus. I finally arrived in Marseille on 1st July 2020, starting the last 18 months of my degree within the Groupe Recherche et Développement (GRD), the instrumentation group at LAM. Hard lockdowns with permanent work-from-home orders were our daily life for months, with much of public life shut down, like in many other parts of the world. Conferences and workshops were held entirely online, work travel for collaborations was limited and only possible nationally. It was during this time that I started holding remote classes on software version control, an effort that was well-received in the French astronomy community. On 15th December 2021, I successfully defended my PhD at the Observatoire de Paris, thus beating the launch of JWST by 10 days.

I will continue my career as a CNES postdoctoral fellow at the Laboratoire d'Etudes Spatiales et d'Instrumentation en Astrophysique (LESIA) in Meudon by Paris. I am looking forward to pursuing my studies of coronagraphy and wavefront sensing and control within the team of the Très haute dynamique (THD) testbed, namely Dr. Johan Mazoyer, Dr. Raphaël Galicher and Dr. Pierre Baudoz.



© Vemaps.com - not to scale

ACKNOWLEDGMENTS

Those who have known me for a long time are aware that I was having an on-and-off inclination to do a PhD for quite a while, with the last word seemingly being spoken at the end of my master's degree, when I decided to take a "real" job instead. Little did I know that the people I met at the place I picked to avoid a PhD are exactly the ones to get me here today. The space on these pages might be sparse, but I want to acknowledge those who have walked this long road with me.

If done right, the process of supervision and mentoring is a two-way street. Both sides, the supervisors and the student must put in their share of energy and time to bring the work to fruition. Over the past several years, I have had the tremendous luck to work with three supervisors, in three different cities, two countries and continents who did exactly that and who enabled me to grow hugely as a researcher: Rémi, Jeff and Laurent.

Rémi, I want to sincerely thank you for the endless support you have given me over the past four years, and for the time you took to share your knowledge with me. By giving me a temporary permanent job at your lab you invited me into a world of research that I never thought I would grow so fond of. You were also the one who first floated the idea of me starting a PhD, and here we are now! Thank you for all the big and small things you taught me while working closely together on a day-to-day basis, inside and outside the lab. I particularly value the moments when you demonstratively sat down on a chair next to HiCAT, put your hand under your chin and said: "You do it now." While the first time this happened, it filled me with mild anxiety, it also showed me you believed I was up to the job. I also need to thank you for giving me the space and urging me to go explore things on my own, and for protecting my time to do so. I am especially thankful that you let me bring in my own ideas into your and the team's work, thus fostering an environment of mutual learning. Most importantly, and I might be speaking for more than myself now, thank you for entering French bulldog mode when one of your own needed support or was facing trouble. The kind of mentoring you gave me is what I hope to pass on to my own mentees in the future. Plus, let's face it, where would I be in life if you hadn't taught me the invaluable skill of sailing a Weta Trimaran?

Jeff, I am grateful how you embraced our initial idea of a joint PhD and went through the trouble of searching and asking around for the funding and collaboration possibilities to make it happen. The way you seemingly easily think of more things to explore and new questions to ask is what I aspire to do with my own work. Thank you for sharing your immense experience in experimental testing and research, and for taking your time to discuss all and any questions I came to you with. Big thanks also for keeping things smooth between the different stakeholders of my PhD, for which I know you did a lot behind the scenes.

Laurent, thank you for joining this journey as my directeur de thèse at a time when you knew neither me nor the Baltimore group personally. You have an extremely instructive teaching style, very pedagogical one might say, and I am grateful that you allowed me to benefit from it. I want to thank you for the respect you brought into our meetings and the many thoughtful insights you led us to with your experience. I also want to state how grateful I am for you getting in touch with me when the world broke down in the initial weeks of the Coronavirus pandemic, to check in on me as I was stuck in a country I was not supposed to be in anymore. I had what I needed to continue life and work safely, but it means a lot when someone actively reaches out.

My PhD journey started at the Makidon lab in Baltimore, where many people have come and gone over the years, but all of them left their mark. I thank all the past and current members of this team for all the dark-hole chasing and Git-hygiene catch-up sessions. Anand, thank you for being my early teacher on everything Fourier optics and the time you set aside to work with me individually. It was a true pleasure to work this closely with an icon in the field, who also never ceases to approach work with such tremendous

joy. Lucie, thank you for teaching me everything about your personal baby, the PASTIS model, as well as the drink, for good measure. Johan, thanks for trying to explain the French higher education system to me (which I still do not understand) and for doing all the hard lifting in the beginning of my PhD by going through the relevant bibliography with me. I am looking forward to working closely with you in the future.

Marshall, thank you for not shying away from taking your time to explain the most basic concepts in optics and software engineering to me, and thank you for just taking a break sometimes, to focus on what really matters. Laurent, thank you for being verbose when others were not, although I still recommend a different naming convention for the variables in your code. You were one of the few to openly acknowledge how detrimental working from home and the pandemic adjustments were and still are for junior people, however necessary these measures were to keep us all healthy and safe. I thank you for helping me navigate the communication with both the wider coronagraphy community as well as the French high-contrast mafia. Scott, thank you for being my only peer PhD student for a year and a half, it was a great pleasure to work alongside you.

Mamadou and Raphaël, I need to thank you for the fantastic weekend tour you gave me of Nice and its surroundings during my long-overdue and Covid-delayed trip to Nice, and for showing me all the exciting work you and your colleagues there are working on.

Chris, thank you for being the most awesome software engineer in astronomy, and a true friend when home is far away. Your patience was crucial in my learning how to code better in Python by reading the official Python documentation from the beginning to the end in one go (who even does that??) and your love for live music is infectious. Keira, I have followed you from the summer school in Toronto over a degree in Leiden to a job in Baltimore, and it is an understatement to say that you paved the way for me. Thank you so much for priming me on the subtleties and not-so-subtleties of American life and for being a pillar of support, both at work and outside. Jules, I want to thank you for your everlasting optimism and undying willingness to sign up for new adventures. Your kind heart and healthy realism managed to pick me up many times and provided lots of crazy and fun moments together.

Ananya, Annie, Emiel, Evelyn, Greg, Heather, Jamie, Julien, Kathryn, Kelsey, Lucas, Maggie, Matt, Meiji, Michele, Nathan, Pete, Rob, Sam, Susan, Théo, Tom, Yinzi, thank you all for turning my time at the lab into a successful and enjoyable one. We will have to catch up on another lab outing to SPiN in Washington, D.C. whenever this hell of a pandemic ebbs away.

My time in Baltimore was marked by many friendly faces, all of whom contributed to me feeling welcome and at home. Bry and Robel, thank you for being valuable friends and a reference against all the confusing things life in America brings sometimes. We never made it to the gym together as originally intended, but our outings to Washington, D.C. and Fells Point, and evenings over Ethiopian and French food were some of my absolute favorite moments we spent together. Katie, thanks a bunch for venturing out with me to neighborhood concerts and events, and for hustling people into table tennis bets they can only lose. John and Sadie, thank you for digesting the initial shock at the onset of the pandemic with me.

This list would be utterly incomplete without the bestest green (both as our shirt color as well as behind the ears) ice hockey team the Baltimore region has seen in recent years, the Wolves. Steve, thank you for always being up to party, even when our hectic lives required quite meticulous planning, and thank you for the endless enthusiasm you brought into all our conversations. Ying, thanks for the company as the only other adult in the very first hockey classes we did together, I am sure we left those kiddos properly traumatized. Damon, thanks for being the biggest fan of my PhD even if its topic is completely removed from what you do in life, the deepest words of support sometimes come from unexpected places. Matt, thanks for being a truly reliable center to us wingers and the enjoyable conversations during our car-sharing commutes to and from the rink. Big thanks to Carlos, Jason, Kerry, Rob, Seth, and all you other crazy hockey nuts I shared the ice with during a pond hockey tournament at -18°C weather, a training camp at the rink of the Pittsburgh Penguins, and on so many other occasions.

After half of my time as a student, I continued work in the sunny south of France. I want to thank the many people at LAM and especially the GRD team, as well as the people who I shared my time in Marseille with. Benoit, thanks for receiving me amid the GRD group as if I had been a member all along. Jorge and Jana, thank you for being my welcoming committee when I first visited Marseille in the late

fall of 2018, showing me around the institute and life in Marseille. Felipe, thank you for shouldering the responsibility of organizing seminar and colloquium talks with me. I want to thank Elodie and Faustine for the wisdom and ideas you shared with me as your junior colleague. Camille, Diana, Gayathri, Ilias, Isabelle, Junais, Lorena, Mahawa, Mathilde, Meriam, Mona, Pia, Sia, thank you and all the other PhD students and colleagues I overlapped with and who were part of our small but strong community trying to survive the academic and pandemic day-to-day.

Alexis, having you become my new office mate right before the big lockdowns that were imposed over the fall, winter and spring of 2020/2021 was probably the best thing that could have happened, and you quickly developed into my closest friend. Simple things like going for (very restricted) walks or cooking dinner together were a much-needed change from the dull, solo confinement days. Towards the end of my PhD, you knew how to get the crew together to have some fun, you always had chocolate ready for the hard days at the office and you could always tell if I needed an ear to rant to or just some quiet to work in peace. Thank you for everything.

Jenna and Max, I thank you for becoming close friends so fast and for your great kindness. Victor, thank you for extending your friendship during a time when you yourself had just made the move back over the pond.

My thanks go to all the members of the DOTA/HRA team at ONERA for constructive conversations and technical exchange. Pablo, it was a great pleasure to start teaching the wonders of Git with an excellent mind like yours and I hope we will not stop here! Alix, thank you for your incessant support and assertion in moments when it showed how hard it still is to be a woman in science. Olivier, you took a huge load off of my shoulders by offering to organize my pot de thèse for me. While the pandemic prevented us from having one (we made up for it with a nice group dinner instead), it shows what a kind character you are, and I want to thank you for that. Antoine, Cyril, Hugo, Hiyam, Léa, Luca, Mehdi, Perrine, Yann, and everybody else on the team, thank you for the company and support during my time as part of the group. I hope to see you guys around after I move to Paris!

Beyond all the colleagues and friends I met during my PhD, I am also indebted to the people who have been part of this journey for longer, in particular the amazing group of friends I found in Leiden. I am forever grateful to the most awesome roommates one can wish for, Andrej and Łukasz. We shared a home, ate and celebrated together, and built friendships that I hope will last a lifetime. Thanks guys!

Patrick, I will never forget how we spent many days and evenings at the various cafés in Leiden, sometimes studying together, sometimes chatting about everything and nothing. Nothing is too hard to do with a friend like you. Just please do not ever sign me up for a 10k again just days before it is happening! Christine, thank you for always being a steady and grounded point of reference while also embracing new adventures and fun times. Eimear, I admire how you always found the right words to explain a situation, especially when I was finding things highly confusing. Thank you for always showing up and for continuously nurturing our friendship over the years.

Bas, Dilovan, Gonzalo, Hiddo, Niloo, Omar, and all the other good people I met at the Leiden Sterrewacht, thanks for a truly impressive time together.

I owe a lot to all of my friends who are scattered around the world, and I want to thank all of you for being at my side for many, many years now - you know who you are. Mia, fate did not play us well as I moved back to Europe just when you moved to the US, but even such obstacles never stopped you from reaching out. Thank you for always having an open ear no matter where in the world you are. Sarah, I think it is safe to say this thesis would not exist without you. Across all cities and countries, you were always there to remind me of what really matters. You showed me that my friends will always be there for me, no matter which path I choose to pursue. For your strong presence and loving support, for your words of encouragement and wake-up calls, I am grateful to you with every fiber of my body. Thank you for never giving up on me.

A special thanks goes out to my core Croatian connection, Petra, Julia and Dora. We have known each other for as long as I can remember, and I am endlessly happy to reconnect with you guys on any possibility we get. Thank you for being there for the past decades. DIP trio for the win!

Finally, none of this would have been possible without the continuous love and support of my family. I

am grateful for the everlasting enthusiasm my late grandfather dedo Drago showed toward my sister and me, excitedly going with us into anything new we were exploring. Thank you for believing in me. I fondly remember my late grandmother baka Danka. You were always hanging on our lips when my sister and I were telling you about our adventures around the world. Thank you for your boundless love. Deep down I know that you are still watching the stars with me.

Thank you to my grandmother baka Rajna, you make me smile every time when you remember and start talking about your past travels to come see me in the Netherlands. My aunt Ivana, thank you for constantly organizing fun and exciting trips with us when we were little, and for teaching us about the world around us. Marta, having you as a younger sister is the biggest gift I have in this world. You taught me so much about compassion and going your own way, you impress me every day. I love you so much.

The biggest heroes of this story are my mom and dad, Ana and Igor. I thank you for your unconditional love, your never ending support and for being the most important anchor in my life. Volim vas do mjeseca i nazad.

